



iSeries

IBM Developer Kit for Java

Version 6





@server

iSeries

IBM Developer Kit for Java

Version 6

Contents

Chapter 1. IBM Developer Kit for Java	1
What's new in V5R2 for IBM Developer Kit for Java	2
Changes to a specific version	3
What's new as of 26 September 2002	3
What's new as of 30 August 2002	3
How to see what's new or changed	3
What's new in V5R2 for Java Development Kit (JDK) 1.1.8	3
What's new in V5R2 for Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.4	3
Print this topic	4
Getting started with the IBM Developer Kit for Java	4
Install the IBM Developer Kit for Java	5
Install a licensed program with the Restore Licensed Program command	5
Support for multiple Java Development Kits	6
Install extensions for the IBM Developer Kit for Java	7
Download and install Java packages on an iSeries server	7
Run your first Hello World Java program	9
Map a network drive to your iSeries server	9
Create a directory on your iSeries server	11
Create a directory using the command entry line	11
Create a directory using iSeries Navigator	12
Create, compile, and run a HelloWorld Java program	12
Create and edit Java source files.	13
With iSeries Access for Windows.	13
On a workstation.	14
With EDTF	14
With Source Entry Utility	14
Work with Java applications using iSeries Navigator.	14
Customize your iSeries server for the IBM Developer Kit for Java.	15
Java classpath	16
Java system properties	18
SystemDefault.properties file	18
Java system properties for Java Development Kit (JDK) 1.1.8	18
os400.stdio.convert and os400.child.stdio.convert system property values	22
os400.stdin, os400.stdout, and os400.stderr system property values	22
os400.verify.checks.disable numeric values	23
Java system properties for Java 2 Software Development Kit (J2SDK), Standard Edition	23
Create an internationalized Java program	27
Time zone environment variable on your iSeries server	28
Configure time zone	28
Java locales	39
Example: Internationalization of dates using the java.util.DateFormat class	41
Example: Internationalization of numeric display using the java.util.NumberFormat class	42
Example: Internationalization of locale-specific data using the java.util.ResourceBundle class	42
Java character encodings	43
File.encoding values and iSeries CCSID	44
Default file.encoding values.	47
Release-to-release compatibility	49
Database access with the IBM Developer Kit for Java	49
Access your iSeries database with the IBM Developer Kit for Java JDBC driver	49
Get started with JDBC.	51
Types of JDBC drivers	51
JDBC requirements.	52
JDBC tutorial	53

Example: JDBC	55
Use JNDI for the examples	59
Connections	59
DriverManager	60
Example: Use native JDBC and Toolbox JDBC concurrently	62
Example: Access property	64
Example: Invalid user ID and password	67
Connection properties	68
Example: Create a UDBDataSource and bind it with JNDI	75
Example: Create a UDBDataSourceBind and set DataSource properties	76
Example: Obtain an initial context before binding UDBDataSource	77
Example: Create a UDBDataSource, and obtain a user ID and password	77
Use DataSources with UDBDataSource	78
DataSource properties.	80
Other DataSource implementations	84
DatabaseMetaData interface for IBM Developer Kit for Java	85
Create a DatabaseMetaData object	85
Retrieve general information	85
Determine feature support	86
Data source limits	86
SQL objects and their attributes	86
Transaction support.	86
Changes in JDBC 3.0	86
Example: DatabaseMetaData interface for IBM Developer Kit for Java	87
Example: Use metadata ResultSets that have more than one column	87
Exceptions	89
SQLException.	90
Example: SQLException	91
SQLWarning	91
DataTruncation	92
Silent truncation	94
Transactions	94
Auto-commit mode	95
Transaction isolation levels	96
Savepoints	97
Distributed transactions	98
Transactions with JTA	98
Use UDBXADatasource support for pooling and distributed transactions	99
XADatasource properties	99
ResultSets and transactions	99
Multiplexing	100
Two-phase commit and transaction logging	100
Example: Use JTA to handle a transaction	101
Example: Multiple connections that work on a transaction	103
Example: Use a connection with multiple transactions	105
Example: Suspended ResultSets	107
Example: End a transaction	109
Example: Suspend and resume a transaction	111
Statement types	114
Statements	115
Example: Use the Statement object's executeUpdate method	116
PreparedStatement	118
Process PreparedStatement.	119
Example: Use PreparedStatement to obtain a ResultSet.	120
Example: ParameterMetaData	122
CallableStatements	123

Process CallableStatements	125
Example: CallableStatement interface for IBM Developer Kit for Java	126
Example: Create a procedure with multiple ResultSets	127
Example: Create a procedure with input and output parameters	128
Example: Create a procedure with return values.	129
ResultSets	131
ResultSet characteristics	131
Example: Sensitive and insensitive ResultSets	133
Example: ResultSet sensitivity	135
Cursor movement	137
Retrieve ResultSet data.	138
Change ResultSets	139
Example: Remove values from a table through another statement's cursor	140
Example: Change values with a statement through another statement's cursor	142
Create ResultSets.	144
Example: ResultSet interface for IBM Developer Kit for Java	144
Example: ResultSetMetaData interface for IBM Developer Kit for Java	145
JDBC object pooling	147
Use DataSource support for object pooling	147
Example: Set up connection pooling with UDBDataSource and UDBConnectionPoolDataSource.	148
Example: Test the performance of connection pooling.	149
ConnectionPoolDataSource properties	149
DataSource-based statement pooling.	151
Example: Test the performance of two DataSources	151
Build your own connection pooling.	152
Batch updates	154
Statement batch update	155
PreparedStatement batch update	156
BatchUpdateException	156
Blocked insert support	156
Advanced data types.	158
Distinct types	158
Large Objects	158
Unsupported SQL3 data types	159
Write code that uses BLOBs	159
Example: BLOB	159
Example: Update BLOBs	161
Example: Use BLOBs	162
Write code that uses CLOBs	163
Example: CLOB	163
Example: Update CLOBs	164
Example: Use CLOBs	165
Write code that uses Datalinks	166
Example: Datalink	166
Example: Distinct types	168
RowSets	169
RowSet characteristics	169
DB2CachedRowSet	170
Use DB2CachedRowSet	171
Create and populate a DB2CachedRowSet	172
Access DB2CachedRowSet data and cursor manipulation	175
Change DB2CachedRowSet data and reflect changes back to the data source	179
Other DB2CachedRowSet features	183
DB2JdbcRowSet	188
DB2JdbcRowSet events	190

Performance tips for the IBM Developer Kit for Java JDBC driver	192
Access databases using IBM Developer Kit for Java DB2 SQLJ support	194
SQLJ tools	195
DB2 SQLJ restrictions	195
Structured Query Language for Java profiles	195
The structured query language for Java (SQLJ) translator (sqlj)	195
Precompile SQL statements in a profile using the DB2 SQLJ profile customizer, db2profc	195
Print the contents of DB2 SQLJ profiles (db2profp and profp)	199
SQLJ profile auditor installer (profdb)	199
Convert a serialized profile instance to Java class format using the SQLJ profile conversion tool (profconv)	200
Embed SQL statements in your Java application	200
Host variables in Structured Query Language for Java	201
Example: Embed SQL Statements in your Java application.	201
Compile and run SQLJ programs	204
Java SQL routines.	205
Use Java SQL routines	206
Java stored procedures	207
JAVA parameter style	207
DB2GENERAL parameter style	209
Restrictions on Java stored procedures	211
Java user-defined scalar functions	211
Parameter style Java.	211
Parameter style DB2GENERAL	212
Restrictions on Java user-defined functions	215
Java user-defined table functions	215
SQLJ procedures that manipulate JAR files	217
SQLJ.INSTALL_JAR	217
SQLJ.REMOVE_JAR	218
SQLJ.REPLACE_JAR	219
SQLJ.UPDATEJARINFO	219
SQLJ.RECOVERJAR	220
Parameter passing conventions for Java stored procedures and UDFs	221
Java with other programming languages	221
Use the Java Native Interface for native methods	222
Java Invocation API	224
Invocation API functions	225
Support for multiple Java virtual machines	226
Example: Java Invocation API	227
Java native methods and threads considerations	232
Native methods and the Java Native Interface	233
Strings in native methods	233
Literal strings in native methods.	234
Convert dynamic strings to and from EBCDIC, Unicode, and UTF-8	234
Examples: Use the Java Native Interface for native methods	235
IBM OS/400 PASE native methods for Java	239
Java OS/400 PASE environment variables	240
Examples: Environment variables for the IBM OS/400 PASE example.	241
Using QIBM_JAVA_PASE_CHILD_STARTUP	241
Managing native method libraries	242
OS/400 PASE and AIX Java Library Naming Conventions	243
Java library search order	243
Java OS/400 PASE error codes.	244
Startup Errors	244
Runtime errors	245
Example: IBM OS/400 PASE native method for Java	245

Run the OS/400 PASE native method for Java example	245
Comparison of Integrated Language Environment and Java	245
Use java.lang.Runtime.exec()	246
Example: Call another Java program with java.lang.Runtime.exec()	247
Example: Call a CL program with java.lang.Runtime.exec().	247
Example: Call a CL command with java.lang.Runtime.exec()	248
Interprocess communications.	249
Use sockets for interprocess communication	249
Example: Use sockets for interprocess communication	249
Use input and output streams for interprocess communication	252
Example: Use input and output streams for interprocess communication	252
Example: Call Java from C	254
Example: Call Java from RPG	254
Java platform	254
Java applets and applications	255
Java virtual machine	255
Java runtime environment	256
Java interpreter.	257
Java JAR and class files	257
Java threads.	257
Sun Microsystems, Inc. Java Development Kit	258
Java packages	259
Java tools.	260
Advanced topics	260
Java classes, packages, and directories.	260
Files in the integrated file system	261
Java file authorities in the integrated file system.	261
Run Java in a batch job	262
Run your Java application on a host that does not have a graphical user interface	263
IBM Developer Kit for Java Remote Abstract Window Toolkit	263
Set up the Remote Abstract Window Toolkit for Java on a remote display	263
Make the Remote Abstract Window Toolkit for Java class files accessible to the remote display	264
Add RAWTGui.zip or RAWTGui.jar to the CLASSPATH of the remote display	265
Start the Remote Abstract Window Toolkit for Java on a remote display	265
Run a Java program using Remote Abstract Window Toolkit	266
Run a Java program using Remote Abstract Window Toolkit with Netscape.	266
Print with the Remote Abstract Window Toolkit	268
Remote Abstract Window Toolkit properties	268
Remote Abstract Window Toolkit SecurityManager restrictions	269
Example: Set up the Remote Abstract Window Toolkit for Java ^(TM) on a Windows remote display	269
Class Broker for Java	269
Set up the Class Broker for Java on a remote display	270
Install Class Broker for Java on an iSeries server	270
Install Class Broker for Java on Windows or UNIX	271
Package content of cbj_1.1.jar	271
Native Abstract Windowing Toolkit	273
Installing NAWT	273
Install OS/400 PASE	273
Install NAWT PTFs	274
Install iSeries Tools for Developers PRPQ	274
Create a VNC password file	274
Configure Java system properties	274
Start the VNC server.	275
Set environment variables	275
Verify installation procedure	275
Installing older versions of iSeries Tools for Developers	275

Determine whether you have the enhanced PRPQ	276
Install VNC	276
Tips on using VNC	276
Start a VNC display server from a CL program	276
End a VNC display server	276
Chapter 2. Java security	279
Java security model	280
Java Cryptography Extension	280
Java Secure Socket Extension	281
Using SSL (JSSE, version 1.0.8)	282
Prepare iSeries server for secure sockets layer support	282
Cryptographic Access Providers.	282
Change your Java code to use socket factories	283
Examples: Change your Java code to use server socket factories	284
Examples: Change your Java code to use client socket factories	286
Change your Java code to use secure sockets layer	287
Examples: Change your Java server to use secure sockets layer	287
Examples: Change your Java client to use secure sockets layer.	289
Select a digital certificate to use	290
Use the digital certificate when you run your Java application.	291
Digital certificates and the -os400.certificateLabel property	291
Digital certificate containers and the -os400.certificateContainer property	291
Using Java Secure Socket Extension, version 1.4	292
Configuring your iSeries server to support JSSE	292
Software requirements	292
Changing JSSE providers	292
Security managers	293
JSSE providers.	293
Pure Java JSSE provider	293
Native iSeries JSSE provider.	293
Changing the default JSSE provider	293
JSSE security properties	293
JSSE Java system properties	295
Properties that work for both providers	295
Properties that work for the iSeries native JSSE provider only	295
Additional information	296
Using the native iSeries JSSE provider	296
Protocol values for the SSLContext.getInstance method	296
Native iSeries KeyStore implementation.	296
Restrictions when using the native iSeries provider	296
Examples: IBM Java Secure Sockets Extension.	297
Example: SSL client using an SSLContext object	298
Example: SSL server using an SSLContext object	299
Chapter 3. Java Authentication and Authorization Service.	303
Prepare and configure an iSeries server for Java Authentication and Authorization Service	303
Java Authentication and Authorization Service samples	305
Chapter 4. IBM Java Generic Security Service (JGSS)	307
JGSS concepts.	308
Principals and credentials	308
Context establishment	309
Message protection and exchange.	309
Resource cleanup and release	309
Security mechanisms	309

Configuring your iSeries server to use IBM JGSS	310
Configuring your iSeries server to use JGSS with J2SDK, version 1.3.	310
Software requirements	310
Configuring your server to use JGSS	310
Changing JGSS providers	310
Security managers	310
Configuring JGSS to use the native iSeries JGSS provider	311
Software requirements	311
Specifying the native iSeries JGSS provider	311
Configuring your iSeries server to use JGSS with J2SDK, version 1.4	312
Changing JGSS providers	312
Security managers	312
JGSS providers.	312
Changing the JGSS provider	312
Using a security manager	313
JVM permissions	313
JAAS permission checks	314
DelegationPermission check	314
ServicePermission check	314
Running IBM JGSS applications	315
Obtaining Kerberos credentials and creating secret keys	315
The Kinit and Ktab tools	315
Using the pure Java JGSS provider	316
Using the native iSeries JGSS provider	316
JAAS Kerberos login interface	316
JAAS and JVM permissions	316
JAAS configuration file options	316
Principal name option	317
Prompting for the principal name and password	317
Credential type option.	318
Configuration and policy files.	318
Kerberos configuration file	318
JAAS configuration file	319
JAAS policy file.	319
Java master security properties file	319
Credentials cache and server key table	320
Developing IBM JGSS applications	320
IBM JGSS application programming steps	321
JGSS transport tokens	321
Sequence of operations in a JGSS application	321
Creating a GSSManager	321
Creating a GSSName	322
Examples: Using GSSName	322
Creating a GSSCredential	322
Creating GSSContext	323
Requesting optional security services.	323
Establishing context	324
Using per-message services	325
Sending messages	325
Receiving messages	326
Deleting context	326
Using JAAS with your JGSS application.	327
Debugging	327
JGSS Debug class	328
Samples: IBM Java Generic Security Service (JGSS).	328
Description of the sample programs	329

Viewing the IBM JGSS samples	329
Viewing the sample programs	329
Viewing the sample configuration and policy files	330
Sample: IBM JGSS non-JAAS client program.	330
Sample: IBM JGSS non-JAAS server program	338
Sample: IBM JGSS JAAS-enabled client program	349
Sample: IBM JGSS JAAS-enabled server program.	351
Sample: Kerberos configuration file	352
Sample: JAAS login configuration file.	353
Sample: JAAS policy file	354
Sample: Java policy file.	355
Samples: Downloading and viewing javadoc information for the IBM JGSS samples	356
Samples: Downloading and running the sample programs	357
Samples: Downloading the IBM JGSS samples	357
Related information	357
Samples: Preparing to run the sample programs	358
Related information	358
Samples: Running the sample programs	358
Related information	359
IBM JGSS javadoc reference information	359
Chapter 5. Tune Java program performance with the IBM Developer Kit for Java.	361
Java runtime performance considerations	361
Cache class loaders	362
Select which mode to use when running a Java program	363
Java interpreter.	365
Static compilation	365
Java static compilation performance considerations	366
Just-In-Time compiler	366
Comparison of Just-In-Time compiler and direct processing	366
Optimization levels	366
Java garbage collection.	367
IBM Developer Kit for Java advanced garbage collection	368
Java garbage collection performance considerations	368
Java Native Method Invocation performance considerations	368
Java method inlining performance considerations	368
Java exception performance considerations	369
Java call trace performance tools	369
Java event trace performance tools	369
Java profiling performance tools	369
Java Virtual Machine Profiler Interface	369
Collect Java performance data	370
Performance Data Collector tool	371
Java Performance Data Converter tool	371
Run the Java Performance Data Converter	372
Example: Run the Java Performance Data Converter.	372
Chapter 6. Commands and tools for the IBM Developer Kit for Java.	373
Java tools that are supported by the IBM Developer Kit for Java.	373
Java tools.	374
Java ajar tool	374
Java appletviewer tool	374
Run the Java appletviewer tool with Remote Abstract Window Toolkit	374
Java extcheck tool	375
Java idlj tool	375
Java jar tool	375

Java jarsigner tool	375
Java javac tool	375
Java javadoc tool	376
Java tools	376
Java javah tool	376
Java javakey tool	377
Java javap tool	377
Java keytool	377
Java native2ascii tool	378
Java policytool	378
Java rmic tool	378
Java rmid tool	378
Java rmiregistry tool	378
Java serialver tool	378
Java tnameserv tool	378
Java command in Qshell	379
CL commands that are supported by Java	380
Analyze Java Virtual Machine (ANZJVM) command	380
Run the ANZJVM command	380
Force a garbage collection cycle	381
Considerations for ANZJVM command	381
Example: ANZJVM command	381
Spooled output file for ANZJVM command	381
Example: Change Java Program (CHGJVAPGM) command	386
Licensed Internal Code option parameter strings	387
Example: Create Java Program (CRTJVAPGM) command	391
Example: Delete Java Program (DLTJVAPGM) command	391
Example: Dump Java Virtual Machine (DMPJVM) command	391
Example: Display Java Program (DSPJVAPGM) command	392
JAVA command	393
Example: Use the Run Java (RUNJVA) command	393
iSeries Navigator commands that are supported by Java	393
Chapter 7. Optional packages	395
Java Naming and Directory Interface	395
IBM JNDI LDAP provider programming guide	396
Create an initial context	397
LDAP V3 URL	398
Server binding and SASL support	399
Search and obtain attributes	401
Add and delete entries in a directory	402
Change attributes	403
Rename a directory entry	403
Referrals and search references	404
LDAP controls	405
Binary attributes	405
Schema	406
SASL plug-in	408
Client-side caching	410
Retrieve the IBMJNDI class version	411
Conformance considerations and additional properties	411
JSSL	412
JavaMail	412
Java Print Service	413
Chapter 8. Debug programs using the IBM Developer Kit for Java	415


Debug a Java program	415
Debug Java programs by using the *DEBUG option	416
Initial debugging displays for Java programs	416
Set breakpoints.	417
Step through Java programs to debug	418
Evaluate variables in Java programs	419
Debug Java and native method programs	419
Debug a Java program from another display	420
QIBM_CHILD_JOB_SNDINQMSG environment variable.	421
Debug Java classes loaded through a custom class loader.	421
Debug servlets	421
Java Platform Debugger Architecture	422
Java Virtual Machine Debug Interface	422
Java Debug Wire Protocol.	422
Start JDWP in QShell	422
Start JDWP from a CL command line.	423
Java Debug Interface	423
Find memory leaks	423
Chapter 9. Troubleshooting the IBM Developer Kit for Java	425
Limitations	425
Find job logs for Java problem analysis	425
Collect data for Java problem analysis	426
Get support for the IBM Developer Kit for Java	427
Chapter 10. Code examples for the IBM Developer Kit for Java.	429
Chapter 11. IBM Developer Kit for Java reference	433
Code disclaimer information	433

Chapter 1. IBM Developer Kit for Java



The IBM(R) Developer Kit for Java^(TM) is optimized for use in an iSeries^(TM) server environment. It uses the compatibility of Java programming and user interfaces, so you can develop your own applications for the iSeries server.

The IBM Developer Kit for Java allows you to create and run Java programs on your iSeries server. The IBM Developer Kit for Java is a compatible implementation of the Sun Microsystems, Inc. Java Technology, so we assume that you are familiar with their Java Development Kit (JDK) documentation. To make it easier for you to work with their information and ours, we provide links to Sun Microsystems, Inc.'s information.

If for any reason our links to Sun Microsystems, Inc. Java Development Kit documentation do not work, refer to their HTML reference documentation for the information that you need. You can find this information on the World Wide Web at The Source for Java Technology java.sun.com .

Select any of these topics for more details on how to use the IBM Developer Kit for Java:

- Print this topic provides details on how to download a printable PDF file or a zipped package of IBM Developer Kit for Java HTML files.
- What's new in V5R2 highlights the latest product and information updates.
- Getting started gives information about installation, configuration, how to create and run simple Hello World Java programs, downloading and installing, and release-to-release compatibility.
- Customization gives instructions on how to customize the configuration of your time zone, system properties, and classpath on your server.
- Compatibility provides information about the compatibility of Java class files from release to release.
- Database access explains how the IBM Developer Kit for Java allows Java programs to access iSeries database files.
- Java with other programming languages shows you how to call code that was written in languages other than Java by using Java Native Interface (JNI), `java.lang.Runtime.exec()`, interprocess communication, and the Java Invocation API.
- Java platform describes the environment for developing and managing Java applets and applications, and consists of the Java language, the Java packages, and the Java virtual machine.
- Advanced topics provides instructions on how to run Java in a batch job and describes the Java file authorities required in the integrated file system to display, run, or debug a Java program.
- Run on a host without a GUI contains information about how to set up and run Java programs with the Remote Abstract Window Toolkit (AWT), Class Broker for Java, or Native Abstract Windowing Toolkit (NAWT).
- Security provides details on adopted authority and explains how you can use SSL to make socket streams secure in your Java application.
- Performance provides information about how to tune your Java performance.
- Commands and tools provides details on how to use the Java commands and Java tools.
- Optional packages lists packages such as JavaMail and Java Naming and Directory Service (JNDI) that you can optionally use to develop your Java applications.
- Debugging explains how to debug Java programs.

- Troubleshooting shows you how to find job logs and collect data for Java program analysis. This topic also provides information about program temporary fixes (PTFs) and getting support for the IBM Developer Kit for Java.
- Code examples links directly to all of the code examples in this information.
- Reference links directly to all of the Javadoc and API reference information.

Note: Read the Code example disclaimer for important legal information.

What's new in V5R2 for IBM Developer Kit for Java

This topic highlights changes to the IBM Developer Kit for Java^(TM) in V5R2. We highlight changes specific to Java Development Kit (JDK) 1.1.8 and Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.4 separately. Updates after the general release of V5R2 appear at the bottom of the following list.

Getting started

- Multiple JDK support has information on each JDK that IBM supports.
- Map a network drive to your iSeries server and Create a directory on your iSeries server have moved to Getting Started.
- Create, compile, and run a Hello World Java program contains several changes.

Customization

- There are new System properties, including properties for Cache class loaders.

Database access

- The JDBC section has had extensive revision.
- The Java stored procedures and Java user-defined scalar functions sections have been added.

Run on a host without a GUI

- See the Run a Java program using Remote Abstract Window Toolkit topic for updates.

Commands and tools

- The Analyze Java Virtual Machine (ANZJVM) command has been added to the CL commands section.
- The Java idlj tool has been added.
- The iSeries Navigator commands that are supported by Java was changed to reflect several changes to iSeries Navigator.

Optional packages

- The JNDI LDAP provider programming guide has been added.

Debugging

- There is new debug support for Java classes loaded through a custom class loader.

Code examples

- More code examples were added.

Print this topic

- Print this topic contains a PDF of the IBM Developer Kit for Java information.

Reference

- The IBM Developer Kit for Java reference section has been added and contains Javadoc and API reference information.

Changes to a specific version

Click on the links below for information specific to the version of your choice:

- Java Development Kit, version 1.1.8
- Java 2 Software Development Kit, Standard Edition, version 1.4

What's new as of 26 September 2002

Support for multiple Java Development Kits

This technical update clarifies the order of precedence that OS/400^(R) uses to determine the default JDK when you have installed more than one JDK.

Calling control language commands from a Java program

This technical update adds information about the delimiters required by different JDKs when calling control language (CL) commands from your Java programs.

What's new as of 30 August 2002

Native Abstract Windowing Toolkit



Changes to Native Abstract Windowing Toolkit (NAWT) information include support for the Java Development Kit version 1.3 and updates to instructions for installing NAWT. The installation instructions now reflect enhancements made to PRPQ 5799-PTL, which is used by NAWT.

Java Generic Security Service

Java Generic Security Service (JGSS) is a new topic that provides information about IBM JGSS, a generic interface for authentication and secure messaging. Under this interface you can plug a variety of security mechanisms based on secret-key, public-key, or other security technologies.

How to see what's new or changed

To help you see where technical changes have been made, this information uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

To find other information about what's new or changed this release, see the Memo to Users .

What's new in V5R2 for Java Development Kit (JDK) 1.1.8

There were no changes in V5R2 specific to the Java^(TM) Development Kit (JDK) 1.1.8.

What's new in V5R2 for Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.4

This topic highlights changes to the IBM Developer Kit for Java^(TM) in V5R2 for the Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.4.

Note: In this section, we only discuss changes that are unique or interesting in the J2SDK, version 1.4. The general information on updates from the What's new for Java Development Kit (JDK) 1.1.8 topic also apply to J2SDK, version 1.4.

Customization

- The time zone environment variable shows unique aspects of setting the time zone variable on an iSeries server. This must be done to use the native `getSystemTimeZoneID()` that is a part of Java 2 SDK (J2SDK), Standard Edition, version 1.3 and higher.

- New system properties for J2SDK were added.

Security

- The Java Authentication and Authorization Service (JAAS) is a standard extension to the Java 2 Software Development Kit, v 1.3 (JDK 1.3) and higher. Currently, Java 2 provides access controls that are based on where the code originated and who signed the code (codesource-based access controls). It lacks, however, the ability to enforce additional access controls based on who runs the code. JAAS provides a framework that augments the Java 2 security model with such support.
- The Java Cryptography Extension (JCE) 1.2 is a standard extension to the Java 2 Software Development Kit (J2SDK), Standard Edition. The JCE implementation on an iSeries server is compatible with the implementation of Sun Microsystems, Inc. This documentation covers the unique aspects of the iSeries implementation. We assume that you are familiar with the general documentation for the JCE extensions.

Optional packages

- The Java Print Service API allows printing on all Java platforms. Java runtime environments and third parties can provide steam generator plugins for producing various formats for printing, such as PDF, Postscript, and AFP.
- The JavaMail API is a set of abstract classes that models an electronic (e-mail) system. The API provides a platform-independent and protocol-independent framework to build Java-based e-mail and messaging applications.

Print this topic


To view or download the PDF version, select IBM Developer Kit for Java^(TM) (about 2159 KB or 450 pages).

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF in your browser (right-click the link above).
2. Click **Save Target As...**
3. Navigate to the directory in which you would like to save the PDF.
4. Click **Save**.

Downloading Adobe Acrobat Reader

If you need Adobe Acrobat Reader to view or print these PDFs, you can download a copy from the Adobe Web site (www.adobe.com/products/acrobat/readstep.html) .

Getting started with the IBM Developer Kit for Java

If you have not yet used the IBM Developer Kit for Java^(TM), follow these steps to install it, configure it, and practice running a simple Hello World Java program.

1. If you are already familiar with the IBM Developer Kit for Java information, see what's new for links to the latest product updates and information.
2. Install the IBM Developer Kit for Java.
3. Configure your system.
4. If you are new to this information and have not yet used the IBM Developer Kit for Java, see Run your first Hello World Java program. This topic illustrates two methods of running a simple Hello World Java program with the IBM Developer Kit for Java. It is a convenient way to see if you have correctly installed the IBM Developer Kit for Java.

5. Now you are ready to create, compile, and run your own Hello World Java program. For how-to steps, see *Create, compile, and run a Hello World Java program*.
6. If you are interested in creating more of your own Java applications, read these topics:
 - *Create and edit Java source files* shows three different ways that you can create and edit your Java source files.
 - *Download and install Java packages on an iSeries server* helps you use Java packages more effectively. It provides details about packages with graphical user interface (GUI), integrated file system and case sensitivity, as well as ZIP file handling and JAR file handling.
 - *Release-to-release compatibility* provides information on compatibility from one release to another.

Install the IBM Developer Kit for Java

Installing the IBM Developer Kit for Java^(TM) allows you to create and run Java programs on your iSeries server.

To install the IBM Developer Kit for Java, perform the following steps:

1. Enter the Go Licensed Program (GO LICPGM) command on the command line.
2. Select option 11 (Install licensed program).
3. Choose option 1 (Install) for licensed program (LP) 5722-JV1 *BASE, and select the option that matches the Java Development Kit (JDK) that you want to install. If the option that you want to install is not displayed in the list, you can add it to the list by entering option 1 (Install) in the option field. Enter 5722JV1 in the licensed program field and your option number in the product option field.

Note: You can install more than one option at once.

Once you have installed the IBM Developer Kit for Java on your iSeries server, you may choose to customize your system.

See *Run your first Hello World Java program* for information on getting started with the IBM Developer Kit for Java.

Install a licensed program with the Restore Licensed Program command

The programs listed in the *Install Licensed Programs* display are those supported by the LICPGM installation when your server was new. Occasionally, new programs become available which are not listed as licensed programs on your server. If this is the case with the program you want to install, you must use the Restore Licensed Program (RSTLICPGM) command to install it.

To install a licensed program with the Restore Licensed Program (RSTLICPGM) command, follow these steps:

1. Put the tape or CD-ROM containing the licensed program in the appropriate drive.
2. On the iSeries command line, type:

RSTLICPGM

and press the **Enter** key.

The *Restore Licensed Program (RSTLICPGM)* display appears.

3. In the *Product* field, type the ID number of the licensed program you want to install.
4. In the *Device* field, specify your install device.

Note: If you are installing from a tape drive, the device ID is usually in the format **TAPXX**, where **XX** is a number, like **01**.

5. Keep the default settings for the other parameters in the *Restore Licensed Program* display. Press the **Enter** key.
6. More parameters appear. Keep these default settings also. Press the **Enter** key. The program begins installing.

When the licensed program is finished installing, the *Restore Licensed Programs* display appears again.

Support for multiple Java Development Kits

Your iSeries server supports multiple Java Development Kits (JDKs) and Java 2 SDK (J2SDK), Standard Edition. Your iSeries server supports the use of multiple JDKs simultaneously, but only through multiple Java virtual machines. A single Java virtual machine runs one specified JDK.

Find the JDK that you are using or want to use, and select the coordinating option to install. You can install more than one JDK at one time. The `java.version` system property determines which JDK to run. Once a Java virtual machine is up and running, changing the `java.version` system property has no effect.

Note: In V5R2, Options 1 (JDK 1.1.6) and 2 (JDK 1.1.7) are no longer available. They cannot be installed or used.

Option	JDK	java.home	java.version
3	1.2	» /QIBM/ProdData/Java400/jdk12/ «	1.2
4	1.1.8	/QIBM/ProdData/Java400/jdk118/	1.1.8
5	1.3	/QIBM/ProdData/Java400/jdk13/	1.3
» 6	1.4	/QIBM/ProdData/Java400/jdk14/	1.4 «

» **Note:** Version 1.3 is the same as Java 2 SDK (J2SDK), Standard Edition, version 1.3. «

For example, these are the results of which options you install and the command you enter.

Install	Enter	Result
Option 3 (1.2)	java Hello	J2SDK, Standard Edition, version 1.2 runs.
Option 4 (1.1.8)	java Hello	JDK 1.1.8 runs, because only one JDK is installed and that is your default.
Option 4 (1.1.8) and Option 3 (1.2)	java Hello	J2SDK, Standard Edition, version 1.2 runs because it is higher.
» All four options installed	java Hello	J2SDK, Standard Edition, version 1.3 runs. «
Option 3 (1.2) and Option 5 (1.3)	java Hello	J2SDK, Standard Edition, version 1.3 runs because it is higher.
Option 4 (1.1.8) and Option 5 (1.3)	java -Djava.version=1.1.8 Hello	JDK 1.1.8 runs, because it was specified.
» Option 5 (1.3) and Option 6 (1.4)	java Hello	J2SDK, Standard Edition, version 1.3 runs. Even though 1.4 is higher in number, 1.3 takes precedence. «

Note: If you install only one JDK, the default JDK is the one you installed. » If you install more than one JDK, the following order of precedence determines the default JDK:

1. Option 5 (1.3)
2. Option 3 (1.2)
3. Option 6 (1.4)
4. Option 4 (1.1.8) «

Install extensions for the IBM Developer Kit for Java

Extensions are packages of Java[™] classes that you can use to extend the functionality of the core platform. Extensions are packaged in one or more ZIP files or JAR files, and are loaded into the Java virtual machine by an extension class loader.

The extension mechanism allows the Java virtual machine to use the extension classes in the same way that the virtual machine uses the system classes. The extension mechanism also provides a way for you to retrieve extensions from specified Uniform Resource Locators (URLs) when they are not already installed in the J2SDK, version 1.2 or higher or Java 2 Runtime Environment, Standard Edition, version 1.2 and higher.

» Some JAR files for extensions are shipped with the iSeries server. « If you would like to install one of these extensions, enter this command:

```
ADDLNK OBJ('/QIBM/ProdData/Java400/ext/extensionToInstall.jar')
NEWLNK('/QIBM/UserData/Java400/ext/extensionToInstall.jar')
LNKTYPE(*SYMBOLIC)
```

Where `extensionToInstall.jar` is the name of the ZIP or JAR file that contains the extension that you want to install.

Note: JAR files of extensions not provided by IBM may be placed in the `/QIBM/UserData/Java400/ext` directory.

When you create a link or add a file to an extension in the `/QIBM/UserData/Java400/ext` directory, the list of files that the extension class loader searches changes for *every Java virtual machine that is running on your iSeries server*. If you do not want to impact the extension class loaders for other Java virtual machines on your iSeries server, but you still want to create a link to an extension or install an extension not shipped by IBM with the iSeries server, follow these steps:

1. Create a directory to install the extensions.
Use either the Make Directory (MKDIR) command from the iSeries command line or the `mkdir` command from the Qshell Interpreter.
2. Place the extension JAR file in the directory created.
3. Add the new directory to the `java.ext.dirs` property.
You can add the new directory to the `java.ext.dirs` property by using the PROP field of the JAVA command from the iSeries command line.

» If the name of your new directory is `/home/username/ext`, the name of your extension file is `extensionToInstall.jar`, and the name of your Java program is `Hello`, then the commands that you enter should look like this:

```
MKDIR DIR('/home/username/ext')

CPY OBJ('/productA/extensionToInstall.jar') TODIR('/home/username/ext') or
copy the file to /home/username/ext using FTP (file transfer protocol).

JAVA Hello PROP((java.ext.dirs '/home/username/ext'))
```

Download and install Java packages on an iSeries server

To download, install, and use Java[™] packages more effectively on an iSeries server, see the following:

- Packages with graphical user interfaces
- Case sensitivity and integrated file system
- ZIP file handling and JAR file handling
- Java extensions framework

Packages with graphical user interfaces

Java programs used with graphical user interface (GUI) require the use of a presentation device with graphical display capabilities. For example, you can use a personal computer, technical workstation, or network computer. The iSeries server provides a Remote Abstract Window Toolkit (AWT) capability. This capability runs applications on an iSeries server by using a full range of graphical capabilities on the appropriate Transmission Control Protocol/Internet Protocol (TCP/IP) attached display device. For more specific installation, setup, and overall use information, see *Set up the Remote Abstract Window Toolkit*.

Case sensitivity and integrated file system

Integrated file system provides file systems, which are both case-sensitive and those that are not with regard to file names. QOpenSys is an example of a case-sensitive file system within the integrated file system. Root, *'/'*, is an example of a case-insensitive file system. For more information about the integrated file system, see the File Systems information in the Integrated file system topic.

➤ Even though a JAR or class may be located in a case-insensitive file system, Java is still a case-sensitive language. While `wrklnk '/home/Hello.class'` and `wrklnk '/home/hello.class'` produce the same results, `JAVA CLASS(Hello)` and `JAVA CLASS(hello)` are calling different classes. ⚡

ZIP file handling and JAR file handling

ZIP files and JAR files contain a set of Java classes. When you use the Create Java Program (CRTJVAPGM) command on one of these files, the classes are verified, converted to an internal machine form, and if specified, transformed to iSeries machine code. You can treat ZIP files and JAR files like any other individual class file. When an internal machine form is associated with one of these files, it remains associated with the file. The internal machine form is used on future runs in place of the class file to improve performance. If you are unsure whether a current Java program is associated with your class file or JAR file, use the Display Java Program (DSPJVAPGM) command to display information about your Java program on your iSeries server.

In previous releases of the IBM Developer Kit for Java, you had to recreate a Java program if you changed the JAR file or ZIP file in any way, because the attached Java program would become unusable. This is no longer true. In many cases, if you change a JAR file or ZIP file, the Java program is still valid and you do not have to recreate it. If partial changes are made, such as when a single class file is updated within a JAR file, you only need to recreate the affected class files that are within the JAR file.

Java programs remain attached to the JAR file after most typical changes to the JAR file. For example, these Java programs remain attached to the JAR file when:

- Changing or recreating a JAR file by using the `ajjar` tool.
- Changing or recreating a JAR file by using the `jar` tool.
- Replacing a JAR file by using the `OS/400 COPY` command or the `Qshell cp` utility.

If you access a JAR file in the integrated file system through iSeries Access for Windows or from a mapped drive on a personal computer (PC), these Java programs remain attached to the JAR file when:

- Dragging and dropping another JAR file onto the existing integrated file system JAR file.
- Changing or recreating the integrated file system JAR file by using the `jar` tool.
- Replacing the integrated file system JAR file by using the `PC copy` command.

When a JAR file is changed or replaced, the Java program that is attached to it is no longer current.

There is one exception in which Java programs do not remain attached to the JAR file. The attached Java programs are destroyed if you use file transfer protocol (FTP) to replace the JAR file. For example, this occurs if you use the `FTP put` command to replace the JAR file.

See Java runtime performance for more detailed information about the performance characteristics of JAR files.

Java extensions framework

In Java 2 SDK, Standard Edition, version 1.2 and higher, extensions are packages of Java classes that you can use to extend the functionality of the core platform. An extension or application is packaged in one or more JAR files. The extension mechanism allows the Java virtual machine to use the extension classes in the same way that the virtual machine uses the system classes. The extension mechanism also provides a way for you to retrieve extensions from specified URLs when they are not already installed in the Java Development Kit (JDK) or Java 2 Runtime Environment, Standard Edition.

See Install extensions for the IBM Developer Kit for Java for information on installing extensions.

Run your first Hello World Java program

You can get your Hello World Java^(TM) program up and running in either of these ways:

1. You can simply run the Hello World Java program that was shipped with the IBM Developer Kit for Java.

To run the program that is included, perform the following steps:

- a. Check that the IBM Developer Kit for Java is installed by entering the Go Licensed Program (GO LICPGM) command. Then, select option 10 (Displayed installed licensed programs). Verify that licensed program 5722-JV1 *BASE and at least one of the options are listed as installed.
 - b. Enter `java Hello` on the iSeries Main Menu command line. Press Enter to run the Hello World Java program.
 - c. If the IBM Developer Kit for Java was installed correctly, Hello World appears in the Java Shell Display. Press F3 (Exit) or F12 (Exit) to return to the command entry display.
 - d. If the Hello World class does not run, check to see that the installation was completed successfully, or see Get support for the IBM Developer Kit for Java for service information.
2. You can also run your own Hello Java program. For more information about how to create your own Hello Java program, see Create, compile, and run a Hello World Java program.

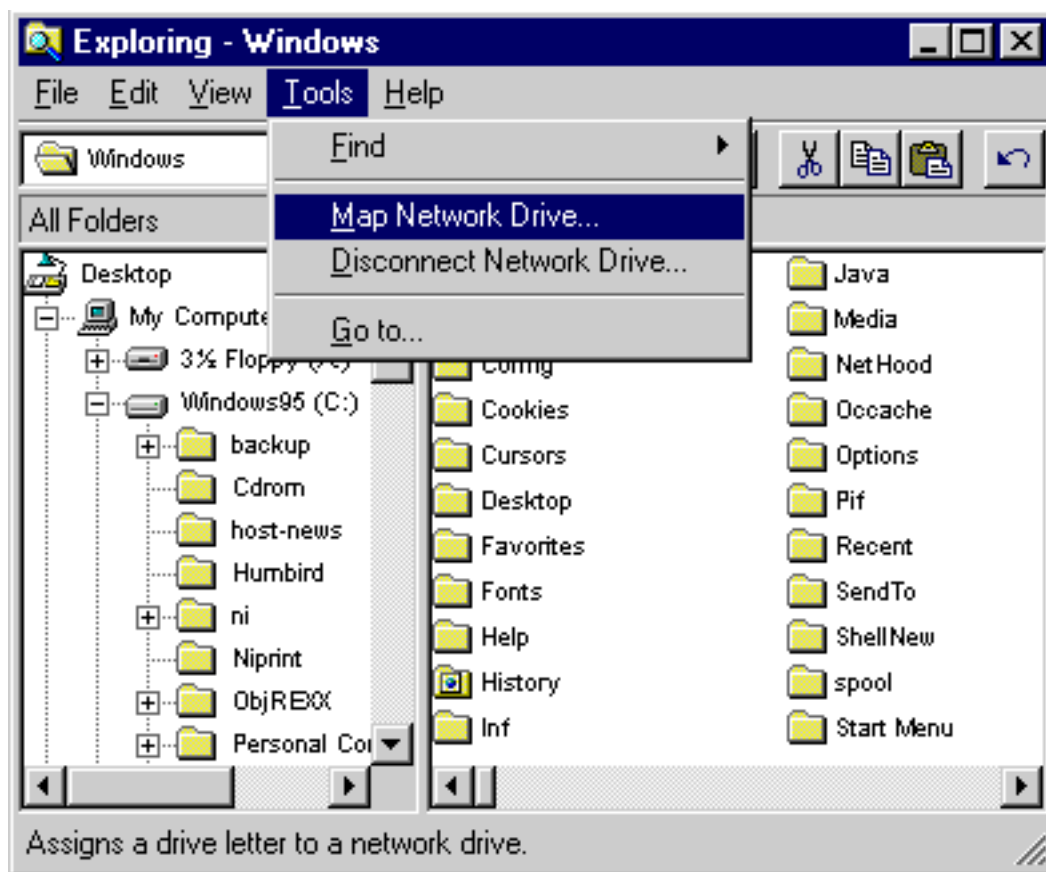
Map a network drive to your iSeries server

» To map a network drive to your iSeries server, make sure that you have iSeries Access for Windows installed on your server and on your workstation. For more information on how to install and configure iSeries Access for Windows, see Installing iSeries Access for Windows.

You must have a connection configured for the iSeries server before you can map a network drive.

To map a network drive, follow these steps:

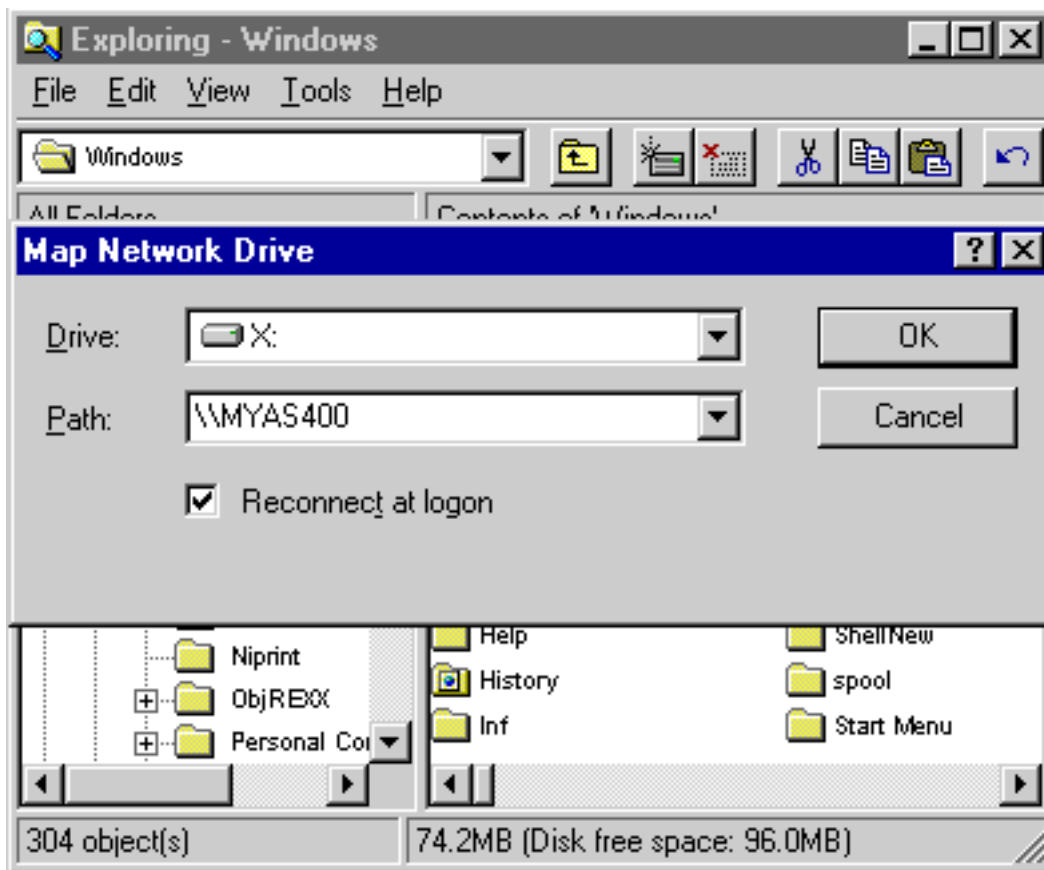
1. Open Windows^(R) Explorer:
 - a. Right-click the **Start** button on your Windows taskbar.
 - b. Click **Explore** in the menu.
2. Select **Map Network Drive** from the **Tools** menu.



3. Select the drive that you want to use to connect to your iSeries server.
4. Type the path name to your server. For example:

\\MYSERVER

where **MYSERVER** is the name of your iSeries server.



5. Check the **Reconnect at logon** box if it is blank.
6. Click **OK** to finish.

Your mapped drive now appears in the **All Folders** section of Windows Explorer. <<

Create a directory on your iSeries server

>> You must create a directory on your iSeries server where you can save your Java[™] applications. There are two ways to do this:

- Create a directory using iSeries Navigator
Choose this option if you have iSeries Access for Windows installed. If you plan to use iSeries Navigator to compile, optimize, and run your Java program, you must select this option to ensure your program is saved in the correct location to perform these operations.
- Create a directory using the command entry line
Choose this option if you do not have iSeries Access for Windows installed.

For information about iSeries Navigator, including installation information, see Getting Started with iSeries Navigator. <<

Create a directory using the command entry line

To create a directory on your iSeries server, follow these steps:

1. Sign on to your iSeries server.
2. On the command line, type:

```
CRTDIR DIR('/mydir')
```

where *mydir* is the name of the directory you are creating.

Press the **Enter** key.

A message appears at the bottom of your screen, stating "**Directory created.**"

Create a directory using iSeries Navigator

» To create a directory on your iSeries server, follow these steps:

1. Open iSeries Navigator.
2. Double-click the name of your server in the **My Connections** window to sign on.
If your server is not listed in the **My Connections** window, follow these steps to add it:
 - a. Click **File** → **Add Connection....**
 - b. Type the name of your server in the **System** field.
 - c. Click **Next**.
 - d. If it is not already entered, enter your User ID in the **Use default user ID, prompt as needed** field.
 - e. Click **Next**.
 - f. Click **Verify Connection**. This confirms that you can connect to the server.
 - g. Click **Finish**.
3. Expand the folder under the connection you want to use. Locate a folder named **File Systems**. If you do not see this folder, the option to install File Systems during the iSeries Navigator installation was not selected. You must install the File Systems option of iSeries Navigator by selecting **Start** → **Programs** → **IBM iSeries Access for Windows** → **Selective Setup**.
4. Expand the **File Systems** folder and locate the **Integrated File System** folder.
5. Expand the **Integrated File System** folder, then expand the **Root** folder. By expanding the **Root** folder, you see the same structure as performing the WRKLNK ('/') command on the iSeries command line.
6. Right-click on the folder where you want to add a subdirectory. Select **New Folder** and enter the name of the subdirectory you want to create.◀◀

Create, compile, and run a HelloWorld Java program


Creating the simple, Hello World Java^(TM) program is a great place to start when becoming familiar with the IBM Developer Kit for Java.

To create, compile, and run your own Hello World Java program, perform the following steps:

1. » Map a network drive to your iSeries server.
2. Create a directory on your iSeries server for your Java applications.◀◀
3. Create the source file as an American Standard Code for Information Interchange (ASCII) text file in the integrated file system. » You can either use an integrated development environment (IDE) product or a text-based editor such as Windows^(R) Notepad to code your Java application. ◀◀
 - a. Name your text file HelloWorld.java. For more information about how you can create and edit your file, see Create and edit Java source files.
 - b. Make sure that your file contains this source code:

```
class HelloWorld {
    public static void main (String args[]) {
        System.out.println("Hello World");
    }
}
```

4. Compile the source file.
 - a. Enter the Work with Environment Variable (WRKENVVAR) command to check the CLASSPATH environment variable. If the CLASSPATH variable does not exist, add it and set it to '.' (the current directory). If the CLASSPATH variable does exist, make sure that the '.' is at the beginning of the path name list. For details about the CLASSPATH environment variable, see Java classpath.




- b. Enter the Start Qshell (STRQSH) command to start the Qshell Interpreter.
 - c. Use the change directory (cd) command to change the current directory to the integrated file system directory that contains the HelloWorld.java file.
 - d. Enter javac followed by the name of the file as you have it saved on your disk. For example, enter javac HelloWorld.java.
5.  Set the file authorities on the class file in the integrated file system.
 6. Optimize the Java application.
 - a. On the *QSH Command Entry* line, type:

```
system "CRTJVAPGM '/mydir/myclass.class' OPTIMIZE(20)"
```


where *mydir* is the path name of the directory in which your Java application is saved, and where *myclass* is the name of your compiled Java application.

Note: You can specify an optimization level of up to 40. An optimization level of 40 increases the efficiency of the Java application, but it also limits debug capabilities. In the early stages of developing a Java application, you may want to set your optimization level at 20 so you can more easily debug your application. See the CRTJVAPGM command and the OPTIMIZE parameter for more information.

- b. Press the **Enter** key.

A message appears, stating that a Java program has been created for your class. 
7. Run the class file.
 - a. Ensure that your Java classpath is set up correctly.
 - b. On the Qshell command line, type java followed by HelloWorld to run your HelloWorld.class with the Java virtual machine. For example, enter java HelloWorld. You can also use the Run Java (RUNJVA) command on your iSeries server to run HelloWorld.class.
 - c. "Hello World" prints to your screen if everything was entered correctly.  The shell prompt (by default, a \$) appears, indicating that the Qshell is ready for another command. 
 - d. Press F3 (Exit) or F12 (Disconnect) to return to the command entry display.



You can also easily compile, optimize, and run your Java application using iSeries Navigator, a graphical user interface for performing tasks on your iSeries server. For instructions, see *Work with Java applications using iSeries Navigator*. For more information on iSeries Navigator, including installation information, see *Getting Started with iSeries Navigator*. 

Create and edit Java source files

You can create and edit Java^(TM) source files in a number of ways:

- With iSeries Access for Windows.
- On a workstation.
- With EDTF.
- With Source Entry Utility.

With iSeries Access for Windows

Java source files are American Standard Code for Information Interchange (ASCII) text files in the integrated file system on iSeries servers.

You can create and edit a Java source file with iSeries Access for Windows and a workstation-based editor.

On a workstation

You can create a Java source file on a workstation. Then, transfer the file to the integrated file system by using file transfer protocol (FTP).

To create and edit Java source files on a workstation:

1. Create the ASCII file on the workstation by using the editor of your choice.
2. Connect to your iSeries server with FTP.
3. Transfer the source file to your directory in the integrated file system as a binary file, so that the file remains in ASCII format.

With EDTF

You can edit files from any file system using the EDTF CL command. It is an editor that is similar to the Source Entry Utility (SEU) for editing stream files or database files. See the EDTF CL command for information.

With Source Entry Utility

You can create a Java source file as a text file by using source entry utility (SEU).

To create a Java source file as a text file by using SEU, perform the following steps:

1. Create a source file member by using SEU.
2. Use the Copy To Stream File (CPYTOSTMF) command to copy the source file member to an integrated file system stream file, while converting the data to ASCII.

If you need to make changes to the source code, change the database member by using SEU and copy the file again.

For information on storing files, see Files in the integrated file system.

Work with Java applications using iSeries Navigator

The iSeries Navigator allows you compile, optimize, and run Java^(TM) applications with point-and-click ease.

Setup requirements

To use iSeries Navigator to work with Java applications, ensure that the following conditions are met:

- iSeries Navigator, which is part of iSeries Access for Windows, must be installed on your workstation. If iSeries Access for Windows is not yet installed, see Getting Started with iSeries Navigator for downloading information.
- You must have your Java application saved to a particular directory on your iSeries server. The correct path for saving your Java application is as follows:

myserver -> File Systems -> Integrated File System -> Root -> home -> mydir

Where **myserver** is the name of your iSeries server, and **mydir** is the name of the directory in which you save your Java applications. For more detailed information on creating a directory in which to save your Java application, see Create a directory using iSeries Navigator.

Compile your Java application using iSeries Navigator



To compile your Java application, follow these steps:

1. Right-click **myfile.java**, where **myfile** is the name of your Java application.
2. Select **Compile Java File**.
3. In the new window, select your JDK version. If you have previously set a classpath according to the instructions given in this guide, it is not necessary to specify a classpath in this window.
4. Click **OK**.

A window opens with a list of errors if any are found in your program. Otherwise, a message stating **1 of 1 Java files were compiled** appears at the bottom of the iSeries Navigator window. A new file called **myfile.class** is created.

Optimize and run your Java application using iSeries Navigator

To optimize and run your Java application, follow these steps:

1. Right-click **myfile.class**.
2.  Select **Associated Java Program**, then click **Run...** to run the Java program. 
3. Click **Advanced**. Select your desired optimization level.

Note: You can specify an optimization level of up to 40. An optimization level of 40 increases the efficiency of the Java application, but it also limits debug capabilities. In the early stages of developing a Java application, you may want to set your optimization level at 20 so you can more easily debug your application. For more information about optimization, see [Optimization levels](#). There is also an option called JIT, which stands for just-in-time compiler. JITs compile code as you need it, making them more efficient than direct processing. For more information about JITs, see [Select which mode to use when running a Java program](#).

4. Click **OK** to close the Advanced Options window.
5. Click **OK** to run the Java program.

Output from the program is displayed in a new window. When the program is finished running, a message stating **Java program completed** appears.

For more information about iSeries Navigator functions, see the **Help** menu on the iSeries Navigator display.

Customize your iSeries server for the IBM Developer Kit for Java

Once you have installed the IBM Developer Kit for Java^(TM) on your iSeries server, you may choose to customize your server.

Time zone configuration

If your Java programs are sensitive to time of day, then time zone configuration may be required.

If the Coordinated Universal Time Offset (QUTCFFSET) system value is set to its default (+00:00), then Java uses the iSeries 400time as the current time. The default value of the user.timezone Java system property setting is UTC.

A QUTCFFSET system value and an updated locale are required for any of these:

- If QUTCFFSET is not set to its default value, thus sensitive to time zone.
- If your Java code expects the user.timezone system property to default to a value other than UTC.
- If you specify the user.timezone Java system property when running the `java` command.

The LC_TOD category of the locale contains the `tname` field, which must be set to the same value that matches your time zone. For details on how to create a locale and format the `tname` field, see [OS/400 globalization](#).

System property configuration

Java system properties determine the environment in which the Java programs run. They are similar to system values or environment variables in OS/400. A number of properties are set when the Java virtual

machine starts. You can choose to use the system property default values or you can specify your own set of default property values by doing either of these steps:

1. Provide a file in /QIBM/UserData/Java400, named SystemDefault.properties. The property values that you specify in this file override the system defaults that IBM provides. This file sets the default system properties for all Java virtual machines that run on your iSeries server.
2. Or, place a SystemDefault.properties file in your own user.home path. This file and the properties that it contains override those in /QIBM/UserData/Java400/SystemDefault.properties. This file sets the default system properties for all Java virtual machines that run for you.

See Run your first Hello World Java program for information on getting started with the IBM Developer Kit for Java.

Java classpath

The Java^(TM) virtual machine uses the Java classpath to find classes during runtime. Java commands and tools also use the classpath to locate classes. The default system classpath, the CLASSPATH environment variable, and the classpath command parameter all determine what directories are searched when looking for a particular class.

Note: In Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.2 and higher, the java.ext.dirs property determines the classpath for the extensions that are loaded. See Install extensions for the IBM Developer Kit for Java for more information.

The default system classpath is system-defined and you cannot change it. On your iSeries server, the default classpath specifies where to find the classes that are part of the IBM Developer Kit, the Remote Abstract Window Toolkit (AWT), and other system classes.

To find any other classes on the system, you must specify the classpath to search by using the CLASSPATH environment variable or the classpath parameter. The classpath parameter that is used on a tool or command overrides the value that is specified in the CLASSPATH environment variable.

You can work with the CLASSPATH environment variable using the Work with Environment Variable (WRKENVVAR) command. From the WRKENVVAR display, you can add or change the CLASSPATH environment variable. The Add Environment Variable (ADDENVVAR) command and Change Environment Variable (CHGENVVAR) command either add or change the CLASSPATH environment variable.

The value of the CLASSPATH environment variable is a list of path names, separated by colons (:), which are searched to find a particular class. A path name is a sequence of zero or more directory names. These directory names are followed by the name of the directory, the ZIP file, or the JAR file that is to be searched in the integrated file system. The components of the path name are separated by the slash (/) character. Use a period (.) to indicate the current working directory.

You can set the CLASSPATH variable in the Qshell environment by using the export utility that is available using the Qshell Interpreter.

These commands add the CLASSPATH variable to your Qshell environment and set it to the value "`./myclasses.zip:/Product/classes.`"

- This command sets the CLASSPATH variable in the Qshell environment:

```
export -s CLASSPATH=./myclasses.zip:/Product/classes
```

- This command sets the CLASSPATH variable from the command line:

```
ADDENVVAR ENVVAR(CLASSPATH) VALUE("./myclasses.zip:/Product/classes")
```

J2SDK searches for classes differently than JDK 1.1.x. JDK 1.1.x searches the system classpath first, then any user-specified classpaths. J2SDK searches the bootstrap classpath first, then the extension directories, then the classpath.

So, the search order for JDK 1.1.x, using the previous example code, is:

1. The default system classpath,
2. The current working directory,
3. The myclasses.zip file that is located in the "root" (/) file system,
4. The classes directory in the Product directory in the "root" (/) file system.

The search order for J2SDK, using the previous example above, is:

1. The bootstrap classpath, which is in the sun.boot.class.path property,
2. The extension directories, which is in the java.ext.dirs property,
3. The current working directory,
4. The myclasses.zip file that is located in the "root" (/) file system,
5. The classes directory in the Product directory in the "root" (/) file system.

When entering the Qshell environment, the CLASSPATH variable is set to the environment variable. The classpath parameter specifies a list of path names. It has the same syntax as the CLASSPATH environment variable. A classpath parameter is available on these tools and commands:

- java command in Qshell
- javac tool
- javah tool
- javap tool
- javadoc tool
- rmic tool
- Run Java (RUNJVA) command

For more information about these commands, see Commands and tools for the IBM Developer Kit for Java. If you use the classpath parameter with any of these command or tools, it ignores the CLASSPATH environment variable.

You can override the CLASSPATH environment variable by using the java.class.path property. You can change the java.class.path property, as well as other properties, by using the SystemDefault.properties file. The values in the SystemDefault.properties files override the CLASSPATH environment variable. For information on the SystemDefault.properties file, see the SystemDefault.properties file.

In JDK 1.1.x, the os400.class.path.system property also affects what is searched when looking for classes. This property can contain one of three values: PRE, POST, or NONE. By default, this property is set to PRE, which causes the default system classpath to be searched before the path. The CLASSPATH environment variable or the classpath parameter specify this path. If you set the os400.class.path.system property to POST, the default system classpath is searched after any user-specified classpath. If the value NONE is used, the default classpath is not searched at all and only the user-specified classpath is searched.

➤ In J2SDK, the -Xbootclasspath option can be used with the same effect. -Xbootclasspath/a: *path* appends *path* to the default bootstrap classpath, /p: *path* prepends *path* to the bootstrap classpath, and /r: *path* replaces the bootstrap classpath with *path*. ⚡

Note: Use caution when specifying NONE or POST because unpredictable results occur if a system class cannot be found or is incorrectly replaced by a user-defined class. Therefore, you should allow the system default classpath to be searched before any user-specified classpath.

See Java system properties for information about how to determine the environment in which Java programs run.

For more information, see the Program and CL Command APIs or the Integrated file system.

Java system properties

Java^(TM) system properties determine the environment in which the Java programs run. They are similar to system values or environment variables in OS/400^(R). A number of properties are set when the Java virtual machine starts.

To view the system properties that this release supports, link to the version of your choice for detailed information:

- Java Development Kit (JDK) 1.1.8
- Java 2 Software Development Kit (J2SDK), Standard Edition [»](#) for versions 1.2, 1.3, and 1.4. [«](#)

SystemDefault.properties file

The SystemDefault.properties file is a standard Java^(TM) properties file. You can specify default properties in a SystemDefault.properties file. The SystemDefault.properties file that is in your home directory takes priority over the SystemDefault.properties file that is in the /QIBM/UserData/Java400 directory. Properties set in the SystemDefault.properties file that is in your home directory are only used for Java virtual machines that you started or users who specify the property user.home = /YourUserHome/.

You can specify the values of properties in the SystemDefault.properties file just like you do in any Java properties file.

Example: SystemDefault.properties file

Note: Read the Code example disclaimer for important legal information.

```
#Comments start with pound sign
#this means always run with JDK 1.3
java.version=1.3
#set my special property
myown.propname=6
```

Java system properties for Java Development Kit (JDK) 1.1.8

Java^(TM) system properties determine the environment in which the Java programs run. They are similar to system values or environment variables in OS/400^(R). A number of properties are set when the Java virtual machine starts.

The system properties are set to these system default values in JDK 1.1.8.

System properties	System default values
awt.toolkit	com.ibm.rawt.client.CToolkit
file.encoding.pkg	sun.io
file.separator	/ (forward slash)
java.class.version	45.3
java.home	See Support for multiple Java Development Kits (JDKs) for details.
java.vendor	IBM Corporation
java.vendor.url	http://www.ibm.com
line.separator	\n
os.arch	PowerPC
os.name	OS/400
os400.class.path.rawt	0

System properties	System default values
os400.class.path.security.check	20 Valid values: 0—No security check 10—equivalent to RUNJVA CHPATH(*IGNORE) 20—equivalent to RUNJVA CHPATH(*WARN) 30—equivalent to RUNJVA CHPATH(*SECURE)
os400.class.path.tools	0
os400.create.type	direct Valid values: interpret—equivalent to RUNJVA OPTIMIZE(*INTERPRET) and INTERPRET(*OPTIMIZE), or INTERPRET(*YES) direct—Otherwise
os400.defineClass.optLevel	20
os400.enbpfrcol	0 Valid values: 0—equivalent to CRTJVAPGM ENBPFRCOL(*NONE) 1—equivalent to CRTJVAPGM ENBPFRCOL(*ENTRYEXIT) 7—equivalent to CRTJVAPGM ENBPFRCOL(*FULL)
os400.interpret	0 Valid values: 0—equivalent to CRTJVAPGM INTERPRET(*NO) 1—equivalent to CRTJVAPGM INTERPRET(*YES)
os400.optimization	10 Valid values: 0—equivalent to CRTJVAPGM OPTIMIZE(*INTERPRET) 10—equivalent to CRTJVAPGM OPTIMIZE(10) 20—equivalent to CRTJVAPGM OPTIMIZE(20) 30—equivalent to CRTJVAPGM OPTIMIZE(30) 40—equivalent to CRTJVAPGM OPTIMIZE(40)
os400.run.mode	program_create_type Valid values: interpret—equivalent to RUNJVA OPTIMIZE(*INTERPRET) and INTERPRET(*OPTIMIZE), or INTERPRET(*YES) program_create_type—Otherwise
os400.stdin.allowed	0
os400.verify.checks.disable	65535 os400.class. This system property value is a string that represents the sum of one or more numeric values. For a list of these values, see os400.verify.checks.disable numeric values.
path.separator	: (colon)

This set of system properties are set based on additional system information.

System properties	Description
file.encoding	Maps the OS/400 job CCSID to the corresponding ISO ASCII CCSID. Also, sets the file.encoding value to the Java value that represents the ISO ASCII coded character set identifier (CCSID). See file.encoding values and iSeries CCSID for a table that shows the relationship between possible file.encoding values and the closest matching iSeries coded character set identifier (CCSID).
java.class.path	Path used to locate classes. Defaults to the default system classpath followed by the user-specified classpath. You can change the java.class.path system property by using the os400.class.path.system system property.
java.compiler	Specifies whether code is compiled with the Just-In-Time (JIT) compiler (jitc) or with both the JIT compiler and direct processing (jitc_de).
java.version	Determines which Java Development Kit (JDK) to run. If you specify a JDK that is not installed, an error message results. If you do not specify a JDK, the most recent JDK is the default. If only one JDK is installed, that is your default. See Support for multiple JDKs for version details.
os.version	Obtains the OS/400 release level from the Retrieve Product Information application program interface (API).
os400.CertificateContainer	Directs Java secure sockets layer (SSL) support to use the specified certificate container for the Java program that was started and the property that was specified. If you specify the os400.secureApplication system property, this system property is ignored. For example, enter -Dos400.certificateContainer=/home/username/mykeyfile.kdb or any other keyfile in the integrated file system.
os400.CertificateLabel	You can specify this system property in conjunction with the os400.CertificateContainer system property. This property lets you select which certificate in the specified container you want secure sockets layer (SSL) to use. For example, enter -Dos400.certificateLabel=myCert, where myCert is the label name that you assign to the certificate through the Digital Certificate Manager (DCM) when you create or import the certificate.
os400.child.stdio.convert	Allows control of the data conversion for stdin, stdout, and stderr in Java. Data conversion occurs by default in the Java virtual machine to convert ASCII data to or from EBCDIC. You can turn these conversions on or off with this property, which affects any child processes that were started by this process by using the runtime.exec() method. See default values.

System properties	Description
os400.class.path.system	<p>PRE (this causes the os400 default system classpath to be prepended onto the user-specified portion of the classpath when constructing the java.class.path system property). Other values are POST (the system default class is appended to the user-specified portion of the classpath) and NONE (only the user-specified classpath is used). The default is PRE.</p> <p>The property is not case sensitive. For example, you can specify NONE, none, noNe. However, the property name is case sensitive. For example, you cannot specify OS400.CLASS.PATH.SYSTEM. To avoid potential problems, you should not change this system property.</p>
<p>» os400.file.create.auth, os400.dir.create.auth</p>	<p>These properties specify authorities assigned to files and directories. Specifying the properties without any values or with unsupported values results in a public authority of *NONE.</p> <p>You can specify os400.file.create.auth=RWX or os400.dir.create.auth=RWX, where R=read, W=write, and X=execute. Any combination of these authorities is valid.</p> <p>«</p>
os400.file.io.mode	<p>Converts the CCSID of the file if it is different than the file.encoding value when you specify TEXT, rather than the default, which is BINARY.</p>
» os400.jit.mmi.threshold	<p>Sets the number of times that a method runs before it is compiled with the JIT Compiler. «</p>
» os400.pool.size	<p>Defines how much space (in kilobytes) to make available for each heap pool in the thread local heap. «</p>
os400.runtime.exec	<ul style="list-style-type: none"> • EXEC (default for 1.3 and higher) - Invoke functions through runtime.exec() using the EXEC interface. This is most compatible with other platforms. • QSHELL (default for 1.2 and lower) - Invoke functions through runtime.exec() using the QSHELL interpreter. This allows the use of variable substitution and invocation of built-in functions.
os400.secureApplication	<p>Associates the Java program that starts when using this system property (os400.secureApplication) with the registered secure application name. You can view registered secure application names by using the Digital Certificate Manager (DCM).</p>
os400.stderr	<p>Allows mapping stderr to a file or socket. See default values.</p>
os400.stdin	<p>Allows mapping stdin to a file or socket. See default values.</p>
os400.stdio.convert	<p>Allows control of the data conversion for stdin, stdout, and stderr in Java. Data conversion occurs by default in the Java virtual machine to convert ASCII data to or from EBCDIC. You can turn these conversions on or off with this property, which affects the current Java program. See default values.</p>
os400.stdout	<p>Allows mapping stdout to a file or socket. See default values.</p>

System properties	Description
user.dir	Current working directory using the getcwd API.
user.home	Retrieves the initial working directory by using the Get API (getpwnam). You can place a SystemDefault.properties file in your user.home path to override the default properties in /QIBM/UserData/Java400/SystemDefault.properties. You can customize the iSeries server to specify your own set of default property values.
user.language	The Java virtual machine uses this system property to read the job LANGID value and uses this value to find the corresponding language.
user.name	The Java virtual machine uses this system property to retrieve the effective user profile name from the Security section (Security.UserName) of the Trusted Computing Base (TCB).
user.region	The Java virtual machine uses this system property to read the job CNTRYID value and uses this value to determine the user region.
user.timezone	The Java virtual machine uses this system property to obtain the time zone name by using the QlgRetrieveLocalInformation API. If no time zone information is available, user.timezone is set to Universal Time Coordinate (UTC).

os400.stdio.convert and os400.child.stdio.convert system property values

The following table shows the system values for both the os400.stdio.convert and os400.child.stdio.convert system properties.

Value	Description
N (default)	No stdio conversion is performed during read or write.
Y	All stdio converts to or from the file.encoding value to job CCSID during read or write.
1	Only stdin data converts from job CCSID to file.encoding during read.
2	Only stdout data converts from file.encoding to job CCSID during write.
3	Both stdin and stdout conversions are performed.
4	Only stderr data converts from file.encoding to job CCSID during write.
5	Both stdin and stderr conversions are performed.
6	Both stdout and stderr conversions are performed.
7	All stdio conversions are performed.

os400.stdin, os400.stdout, and os400.stderr system property values

The following table shows the system values for os400.stdin, os400.stdout, and os400.stderr system properties.

Value	Example name	Description	Example
File	SomeFileName	SomeFileName is an absolute path or a path relative to the current directory.	file:/QIBM/UserData/Java400/Output.file
Port	HostName	Port address	port:myhost:2000
Port	TCPAddress	Port address	port:1.1.11.111:2000

os400.verify.checks.disable numeric values

The os400.verify.checks.disable system property value is a string that represents the sum of one or more numeric values from this list:

Value	Description
1	Bypass access checks for local classes: Indicates that you want the Java ^(TM) virtual machine to bypass access checks on private and protected fields and methods for classes that are loaded from the local file system. It is helpful when transferring applications, which contain inner classes that refer to private and protected methods and fields of their enclosing classes.
2	Suppress NoClassDefFoundError during early load: Indicates that you want the Java virtual machine to ignore NoClassDefFoundErrors, which occur during early verification checks for typecasting and field or method access.
4	Allow LocalVariableTable checking to be bypassed: Indicates that if you encounter an error in the LocalVariableTable of a class, the class operates as if the LocalVariableTable does not exist. Otherwise errors in the LocalVariableTable result in a ClassFormatError.
7	Value used at runtime.

You can indicate the value in decimal, hexadecimal, or octal format. It ignores values that are less than zero. For example, to select the first two values from the list, use this iSeries command syntax:

```
JAVA CLASS>Hello) PROP((os400.verify.checks.disable 3))
```

Java system properties for Java 2 Software Development Kit (J2SDK), Standard Edition

Java^(TM) system properties determine the environment in which the Java programs run. They are similar to system values or environment variables in OS/400. A number of properties are set when the Java virtual machine starts.

» The system properties are set to these system default values in Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.4. A number of system properties have different default values if you go through the JAVA or RUNJAVA CL command than if you use the Java Native Interface (JNI) Invocation API. The following table reflects the use of the API. «

System properties	System default values
awt.toolkit	For JDK 1.1.x, the default value is com.ibm.rawt.client.CToolkit. For J2SDK, the default value is com.ibm.rawt2.ahost.java.awt.AHToolkit.
file.encoding.pkg	sun.io
file.separator	/ (forward slash)
java.class.version	» 48.0 «
java.ext.dirs	» /QIBM/ProdData/Java400/jdk14/lib/ext:/QIBM/UserData/Java400/ext «
java.home	See Support for multiple Java Development Kits (JDKs) for details.
java.library.path	OS/400 library list
java.policy	» /QIBM/ProdData/Java400/jdk14/lib/security/java.policy «
java.specification.name	Java Language Specification
java.specification.vendor	Sun Microsystems, Inc.


System properties	System default values
java.specification.version	» 1.4 «
» sun.boot.class.path «	Class_Path_Sys
java.use.policy	true
java.vendor	IBM Corporation
java.vendor.url	http://www.ibm.com
java.vm.name	» Classic VM «
java.vm.specification.name	Java Virtual Machine Specification
java.vm.specification.vendor	Sun Microsystems, Inc.
java.vm.specification.version	» 1.0 «
java.vm.vendor	IBM Corporation
java.vm.version	» 1.4 «
line.separator	\n
os.arch	PowerPC
os.name	OS/400
os400.class.path.rawt	0
os400.class.path.security.check	20 Valid values: 0—No security check 10—equivalent to RUNJVA CHPATH(*IGNORE) 20—equivalent to RUNJVA CHPATH(*WARN) 30—equivalent to RUNJVA CHPATH(*SECURE)
os400.class.path.tools	0
os400.create.type	» interpret « Valid values: interpret—equivalent to RUNJVA OPTIMIZE(*INTERPRET) and INTERPRET(*OPTIMIZE), or INTERPRET(*YES) direct—Otherwise
os400.defineClass.optLevel	20
os400.enbpfrcol	0 Valid values: 0—equivalent to CRTJVAPGM ENBPFRCOL(*NONE) 1—equivalent to CRTJVAPGM ENBPFRCOL(*ENTRYEXIT) 7—equivalent to CRTJVAPGM ENBPFRCOL(*FULL) For a nonzero value, the JIT generates *JVAENTRY, *JVAEXIT, *JVAPRECALL and *JVAPOSTCALL events.
os400.interpret	0 Valid values: 0—equivalent to CRTJVAPGM INTERPRET(*NO) 1—equivalent to CRTJVAPGM INTERPRET(*YES)

System properties	System default values
os400.optimization	>> 0 << Valid values: 0—equivalent to CRTJVAPGM OPTIMIZE(*INTERPRET) 10—equivalent to CRTJVAPGM OPTIMIZE(10) 20—equivalent to CRTJVAPGM OPTIMIZE(20) 30—equivalent to CRTJVAPGM OPTIMIZE(30) 40—equivalent to CRTJVAPGM OPTIMIZE(40)
os400.run.mode	>> jitc_de << Valid values: interpret—equivalent to RUNJVA OPTIMIZE(*INTERPRET) and INTERPRET(*OPTIMIZE), or INTERPRET(*YES) program_create_type jitc_de—Otherwise
os400.stdin.allowed	0
os400.verify.checks.disable	65535 This system property value is a string that represents the sum of one or more numeric values. For a list of these values, see os400.verify.checks.disable numeric values.
path.separator	: (colon)

This set of system properties are set based on additional system information.

System properties	Description
file.encoding	Maps the OS/400 job CCSID to the corresponding ISO ASCII CCSID. Also, sets the file.encoding value to the Java value that represents the ISO ASCII coded character set identifier (CCSID). See file.encoding values and iSeries CCSID for a table that shows the relationship between possible file.encoding values and the closest matching iSeries coded character set identifier (CCSID).
java.class.path	Path used to locate classes. Defaults to the user-specified classpath.
java.compiler	Specifies whether code is compiled with the Just-In-Time (JIT) compiler (jitic) or with both the JIT compiler and direct processing (jitic_de).
java.version	Determines which Java Development Kit (JDK) to run. If you specify a JDK that is not installed, an error message results. If you do not specify a JDK, the most recent JDK is the default. If only one JDK is installed, that is your default. See Support for multiple JDKs for version details.
os.version	Obtains the OS/400 release level from the Retrieve Product Information application program interface (API).
os400.CertificateContainer	Directs Java secure sockets layer (SSL) support to use the specified certificate container for the Java program that was started and the property that was specified. If you specify the os400.secureApplication system property, this system property is ignored. For example, enter -Dos400.certificateContainer=/home/username/mykeyfile.kdb or any other keyfile in the integrated file system.

System properties	Description
os400.CertificateLabel	You can specify this system property in conjunction with the os400.CertificateContainer system property. This property lets you select which certificate in the specified container you want secure sockets layer (SSL) to use. For example, enter-Dos400.certificateLabel=myCert, where myCert is the label name that you assign to the certificate through the Digital Certificate Manager (DCM) when you create or import the certificate.
os400.child.stdio.convert	Allows control of the data conversion for stdin, stdout, and stderr in Java. Data conversion occurs by default in the Java virtual machine to convert ASCII data to or from EBCDIC. You can turn these conversions on or off with this property, which affects any child processes that were started by this process by using the runtime.exec() method. See default values.
os400.class.path.system	» This system property is ignored for J2SDK. «
» os400.define.class.cache.file	This property specifies the name of a JAR or ZIP file. The default is null. See Cache class loaders. «
» os400.define.class.cache.hours	This property is a decimal value. The default value is 168 and the maximum decimal value is 9999. See Cache class loaders. «
» os400.define.class.cache.maxpgms	This property is a decimal value. The default value is 5000 and the maximum decimal value is 40000. See Cache class loaders. «
os400.exception.trace	Specifying this property causes the most recent exceptions to be sent to standard output when the JVM exits. The value assigned to this property is currently ignored, but that may change in the future. This property is used purely for debugging.
» os400.file.create.auth, os400.dir.create.auth	These properties specify authorities assigned to files and directories. Specifying the properties without any values or with unsupported values results in a public authority of *NONE. You can specify os400.file.create.auth=RWX or os400.dir.create.auth=RWX, where R=read, W=write, and X=execute. Any combination of these authorities is valid. «
os400.file.io.mode	Converts the CCSID of the file if it is different than the file.encoding value when you specify TEXT, rather than the default, which is BINARY.
» os400.jit.mmi.threshold	Sets the number of times that a method runs before it is compiled with the JIT Compiler. «
» os400.pool.size	Defines how much space (in kilobytes) to make available for each heap pool in the thread local heap. «
os400.runtime.exec	<ul style="list-style-type: none"> EXEC (default for 1.3 and higher) - Invoke functions through runtime.exec() using the EXEC interface. This is most compatible with other platforms. QSHHELL (default for 1.2 and lower) - Invoke functions through runtime.exec() using the QSHHELL interpreter. This allows the use of variable substitution and invocation of built-in functions.

System properties	Description
os400.secureApplication	Associates the Java program that starts when using this system property (os400.secureApplication) with the registered secure application name. You can view registered secure application names by using the Digital Certificate Manager (DCM).
os.400.security.properties	Allows full control over which java.security file you use. When you specify this property, J2SDK will not use any other java.security files, including the J2SDK specific java.security default. 
os400.stderr	Allows mapping stderr to a file or socket. See default values.
os400.stdin	Allows mapping stdin to a file or socket. See default values.
os400.stdin.allowed	Specifies whether stdin is allowed (1) or not allowed (0). If the caller is running a batch job, stdin should not be allowed. The default 0.
os400.stdio.convert	Allows control of the data conversion for stdin, stdout, and stderr in Java. Data conversion occurs by default in the Java virtual machine to convert ASCII data to or from EBCDIC. You can turn these conversions on or off with this property, which affects the current Java program. See default values.
os400.stdout	Allows mapping stdout to a file or socket. See default values.
user.dir	Current working directory using the getcwd API.
user.home	Retrieves the initial working directory by using the Get API (getpwnam). You can place a SystemDefault.properties file in your user.home path to override the default properties in /QIBM/UserData/Java400/SystemDefault.properties. You can customize the iSeries server to specify your own set of default property values.
user.language	The Java virtual machine uses this system property to read the job LANGID value and uses this value to find the corresponding language.
user.name	The Java virtual machine uses this system property to retrieve the effective user profile name from the Security section (Security.UserName) of the Trusted Computing Base (TCB).
user.region	The Java virtual machine uses this system property to read the job CNTRYID value and uses this value to determine the user region.
user.timezone	The Java virtual machine uses this system property to obtain the time zone name by using the QlgRetrieveLocalInformation API. If no time zone information is available, user.timezone is set to Universal Time Coordinate (UTC).

Create an internationalized Java program

If you need to customize a Java^(TM) program for a specific region of the world, you can create an internationalized Java program with Java locales.

To create an internationalized Java program, perform the following steps:

1. Isolate the locale-sensitive code and data. For example, strings, dates, and numbers in your program.
2. Set or get the locale using the `Locale` class.
3. Format dates and numbers to specify a locale if you do not want to use the default locale.
4. Create resource bundles to handle strings and other locale-sensitive data.

To carry out these tasks in your own Java programs, refer to these examples:

- Internationalization of dates using the `java.util.DateFormat` class
- Internationalization of numeric display using the `java.util.NumberFormat` class
- Internationalization of locale-specific data using the `java.util.ResourceBundle` class

For more information about internationalization, click on any of these links:

- OS/400 globalization
- Internationalization by Sun Microsystems, Inc.

Time zone environment variable on your iSeries server

With Java^(TM) 2 SDK (J2SDK), Standard Edition, version 1.4, you can set your Java Virtual Machine's time zone by using the native method `getSystemTimeZoneID()`. The iSeries server uses the `*LOCALE` object that is part of the `*ENV` object. Set the `tname` field in the `*LOCALE` object to the appropriate system value. That value is then returned as the associated Java string object from `getSystemTimeZoneID()`.

Configure time zone: The JVM requires both the `QUTCOffset` system value and time of day information in the `LOCALE` for the current job to be set to determine local time correctly. `QUTCOffset` is a system value specifying the number of hours difference between the current local time and the universal time coordinate (UTC). For central standard time (CST), this would be -6:00. For central daylight time (CDT), the correct value is -5:00. The `QUTCOffset` value allows the JVM to determine the correct value for UTC.

`LOCALE` information for a job is set by creating a `*LOCALE` object containing time of day information and specifying that `*LOCALE` for the job using the `QLOCALE` system value or the `LOCALE` keyword on the user profile for the job. Details on creating and using `LOCALE` can be found in the OS/400 globalization publication.

Setting the `*LOCALE` information correctly allows the JVM to default the `user.timezone` property to the correct time zone. The `user.timezone` property can be set manually on the command line to override the default setting provided by the `*LOCALE` object.

The following is an example of the `LC_TOD` information that must be included in the `*LOCALE` object to configure the correct time zone for Java:

```
LC_TOD

% TZDIFF is number of minutes difference from GMT
tzdiff -300
% Timezone name (this is the value that you would have passed to
% the JVM as the user.timezone property.) See abbreviations later
% in this document.
tname "<C><S><T>"
% Name used for daylight savings time.
dstname "<C><D><T>"
% DST Start in this part of the US is the first Sunday in April at 2am
dststart 4,1,1,7200
% DST End in this area of US is Last Sunday in October.
dstend 10,-1,1,7200
% shift in seconds
dstshift 3600

END LC_TOD
```

The following table indicates the system values and the associated Java string objects.

Note: The system value "Hong Kong" represents China (Hong Kong S.A.R.).

System value	Java string object
Africa/Abidjan	Africa/Abidjan
Africa/Accra	Africa/Accra
Africa/Addis_Ababa	Africa/Addis_Ababa
Africa/Algiers	Africa/Algiers
Africa/Asmera	Africa/Asmera
Africa/Bamako	GMT
Africa/Bangui	Africa/Bangui
Africa/Banjul	Africa/Banjul
Africa/Bissau	Africa/Bissau
Africa/Blantyre	Africa/Blantyre
Africa/Brazzaville	Africa/Luanda
Africa/Bujumbura	Africa/Bujumbura
Africa/Cairo	Africa/Cairo
Africa/Casablanca	Africa/Casablanca
Africa/Ceuta	Europe/Paris
Africa/Conakry	Africa/Conakry
Africa/Dakar	Africa/Dakar
Africa/Dar_es_Salaam	Africa/Dar_es_Salaam
Africa/Djibouti	Africa/Djibouti
Africa/Douala	Africa/Douala
Africa/El_Aaiun	Africa/Casablanca
Africa/Freetown	Africa/Freetown
Africa/Gaborone	Africa/Gaborone
Africa/Harare	Africa/Harare
Africa/Johannesburg	Africa/Johannesburg
Africa/Kampala	Africa/Kampala
Africa/Khartoum	Africa/Khartoum
Africa/Kigali	Africa/Kigali
Africa/Kinshasa	Africa/Kinshasa
Africa/Lagos	Africa/Lagos
Africa/Libreville	Africa/Libreville
Africa/Lome	Africa/Lome
Africa/Luanda	Africa/Luanda
Africa/Lubumbashi	Africa/Lubumbashi
Africa/Lusaka	Africa/Lusaka
Africa/Malabo	Africa/Malabo
Africa/Maputo	Africa/Maputo
Africa/Maseru	Africa/Maseru

System value	Java string object
Africa/Mbabane	Africa/Mbabane
Africa/Mogadishu	Africa/Mogadishu
Africa/Monrovia	Africa/Monrovia
Africa/Nairobi	Africa/Nairobi
Africa/Ndjamena	Africa/Ndjamena
Africa/Niamey	Africa/Niamey
Africa/Nouakchott	Africa/Nouakchott
Africa/Ouagadougou	Africa/Ouagadougou
Africa/Porto-Novo	Africa/Porto-Novo
Africa/Sao_Tome	Africa/Sao_Tome
Africa/Timbuktu	Africa/Timbuktu
Africa/Tripoli	Africa/Tripoli
Africa/Tunis	Africa/Tunis
Africa/Windhoek	Africa/Windhoek
America/Adak	America/Adak
America/Anchorage	America/Anchorage
America/Anguilla	America/Anguilla
America/Antigua	America/Antigua
America/Araguaina	America/Sao_Paulo
America/Aruba	America/Aruba
America/Asuncion	America/Asuncion
America/Atka	America/Adak
America/Barbados	America/Barbados
America/Belize	America/Belize
America/Bogota	America/Bogota
America/Boise	America/Denver
America/Buenos_Aires	America/Buenos_Aires
America/Cancun	America/Chicago
America/Caracas	America/Caracas
America/Cayenne	America/Cayenne
America/Cayman	America/Cayman
America/Chicago	America/Chicago
America/Chihuahua	America/Denver
America/Costa_Rica	America/Costa_Rica
America/Cuiaba	America/Cuiaba
America/Curacao	America/Curacao
America/Dawson	America/Los_Angeles
America/Dawson_Creek	America/Dawson_Creek
America/Denver	America/Denver
America/Detroit	America/New_York
America/Dominica	America/Dominica

System value	Java string object
America/Edmonton	America/Edmonton
America/EI_Salvador	America/EI_Salvador
America/Ensenada	America/Los_Angeles
America/Fort_Wayne	America/Indianapolis
America/Fortaleza	America/Fortaleza
America/Glace_Bay	America/Halifax
America/Godthab	America/Godthab
America/Goose_Bay	America/Thule
America/Grand_Turk	America/Grand_Turk
America/Grenada	America/Grenada
America/Guadeloupe	America/Guadeloupe
America/Guatemala	America/Guatemala
America/Guayaquil	America/Guayaquil
America/Guyana	America/Guyana
America/Halifax	America/Halifax
America/Havana	America/Havana
America/Indiana/Indianapolis	America/Indianapolis
America/Indianapolis	America/Indianapolis
America/Inuvik	America/Denver
America/Iqaluit	America/New_York
America/Jamaica	America/Jamaica
America/Juneau	America/Anchorage
America/La_Paz	America/La_Paz
America/Lima	America/Lima
America/Los_Angeles	America/Los_Angeles
America/Louisville	America/New_York
America/Managua	America/Managua
America/Manaus	America/Manaus
America/Martinique	America/Martinique
America/Mazatlan	America/Mazatlan
America/Menominee	America/Winnipeg
America/Mexico_City	America/Mexico_City
America/Miquelon	America/Miquelon
America/Montevideo	America/Montevideo
America/Montreal	America/Montreal
America/Montserrat	America/Montserrat
America/Nassau	America/Nassau
America/New_York	America/New_York
America/Nipigon	America/New_York
America/Nome	America/Anchorage
America/Noronha	America/Noronha

System value	Java string object
America/Panama	America/Panama
America/Pangnirtung	America/Thule
America/Paramaribo	America/Paramaribo
America/Phoenix	America/Phoenix
America/Port-au-Prince	America/Port-au-Prince
America/Port_of_Spain	America/Port_of_Spain
America/Porto_Acre	America/Porto_Acre
America/Puerto_Rico	America/Puerto_Rico
America/Rainy_River	America/Chicago
America/Rankin_Inlet	America/Chicago
America/Regina	America/Regina
America/Santiago	America/Santiago
America/Santo_Domingo	America/Santo_Domingo
America/Sao_Paulo	America/Sao_Paulo
America/Scoresbysund	America/Scoresbysund
America/Shiprock	America/Denver
America/St_Johns	America/St_Johns
America/St_Kitts	America/St_Kitts
America/St_Lucia	America/St_Lucia
America/St_Thomas	America/St_Thomas
America/St_Vincent	America/St_Vincent
America/Tegucigalpa	America/Tegucigalpa
America/Thule	America/Thule
America/Thunder_Bay	America/New_York
America/Tijuana	America/Tijuana
America/Tortola	America/Tortola
America/Vancouver	America/Vancouver
America/Virgin	America/St_Thomas
America/Whitehorse	America/Los_Angeles
America/Winnipeg	America/Winnipeg
America/Yakutat	America/Anchorage
America/Yellowknife	America/Denver
Antarctica/Casey	Antarctica/Casey
Antarctica/DumontDUrville	Antarctica/DumontDUrville
Antarctica/Mawson	Antarctica/Mawson
Antarctica/McMurdo	Antarctica/McMurdo
Antarctica/Palmer	Antarctica/Palmer
Antarctica/South_Pole	Antarctica/McMurdo
Arctic/Longyearbyen	Europe/Oslo
Asia/Aden	Asia/Aden
Asia/Almaty	Asia/Almaty

System value	Java string object
Asia/Amman	Asia/Amman
Asia/Anadyr	Asia/Anadyr
Asia/Aqtau	Asia/Aqtau
Asia/Aqtobe	Asia/Aqtobe
Asia/Ashkhabad	Asia/Ashkhabad
Asia/Baghdad	Asia/Baghdad
Asia/Bahrain	Asia/Bahrain
Asia/Baku	Asia/Baku
Asia/Bangkok	Asia/Bangkok
Asia/Beirut	Asia/Beirut
Asia/Bishkek	Asia/Bishkek
Asia/Brunei	Asia/Brunei
Asia/Calcutta	Asia/Calcutta
Asia/Chungking	Asia/Shanghai
Asia/Colombo	Asia/Colombo
Asia/Dacca	Asia/Dacca
Asia/Damascus	Asia/Damascus
Asia/Dubai	Asia/Dubai
Asia/Dushanbe	Asia/Dushanbe
Asia/Gaza	Asia/Amman
Asia/Harbin	Asia/Shanghai
Asia/Hong_Kong	Asia/Hong_Kong
Asia/Irkutsk	Asia/Irkutsk
Asia/Istanbul	Europe/Istanbul
Asia/Jakarta	Asia/Jakarta
Asia/Jayapura	Asia/Jayapura
Asia/Jerusalem	Asia/Jerusalem
Asia/Kabul	Asia/Kabul
Asia/Kamchatka	Asia/Kamchatka
Asia/Karachi	Asia/Karachi
Asia/Kashgar	Asia/Shanghai
Asia/Katmandu	Asia/Katmandu
Asia/Krasnoyarsk	Asia/Krasnoyarsk
Asia/Kuala_Lumpur	Asia/Kuala_Lumpur
Asia/Kuwait	Asia/Kuwait
Asia/Macao	Asia/Macao
Asia/Magadan	Asia/Magadan
Asia/Manila	Asia/Manila
Asia/Muscat	Asia/Muscat
Asia/Nicosia	Asia/Nicosia
Asia/Novosibirsk	Asia/Novosibirsk

System value	Java string object
Asia/Omsk	Asia/Novosibirsk
Asia/Phnom_Penh	Asia/Phnom_Penh
Asia/Pyongyang	Asia/Pyongyang
Asia/Qatar	Asia/Qatar
Asia/Rangoon	Asia/Rangoon
Asia/Riyadh	Asia/Riyadh
Asia/Saigon	Asia/Saigon
Asia/Seoul	Asia/Seoul
Asia/Shanghai	Asia/Shanghai
Asia/Singapore	Asia/Singapore
Asia/Taipei	Asia/Taipei
Asia/Tashkent	Asia/Tashkent
Asia/Tbilisi	Asia/Tbilisi
Asia/Tehran	Asia/Tehran
Asia/Tel_Aviv	Asia/Jerusalem
Asia/Thimbu	Asia/Thimbu
Asia/Tokyo	Asia/Tokyo
Asia/Ujung_Pandang	Asia/Ujung_Pandang
Asia/Ulan_Bator	Asia/Ulan_Bator
Asia/Urumqi	Asia/Shanghai
Asia/Vientiane	Asia/Vientiane
Asia/Vladivostok	Asia/Vladivostok
Asia/Yakutsk	Asia/Yakutsk
Asia/Yekaterinburg	Asia/Yekaterinburg
Asia/Yerevan	Asia/Yerevan
Atlantic/Azores	Atlantic/Azores
Atlantic/Bermuda	Atlantic/Bermuda
Atlantic/Canary	Atlantic/Canary
Atlantic/Cape_Verde	Atlantic/Cape_Verde
Atlantic/Faeroe	Atlantic/Faeroe
Atlantic/Jan_Mayen	Atlantic/Jan_Mayen
Atlantic/Madeira	Europe/London
Atlantic/Reykjavik	Atlantic/Reykjavik
Atlantic/South_Georgia	Atlantic/South_Georgia
Atlantic/St_Helena	Atlantic/St_Helena
Atlantic/Stanley	Atlantic/Stanley
Australia/ACT	Australia/Sydney
Australia/Adelaide	Australia/Adelaide
Australia/Brisbane	Australia/Brisbane
Australia/Broken_Hill	Australia/Broken_Hill
Australia/Canberra	Australia/Sydney

System value	Java string object
Australia/Darwin	Australia/Darwin
Australia/Hobart	Australia/Hobart
Australia/LHI	Australia/Lord_Howe
Australia/Lord_Howe	Australia/Lord_Howe
Australia/Melbourne	Australia/Sydney
Australia/NSW	Australia/Sydney
Australia/North	Australia/Darwin
Australia/Perth	Australia/Perth
Australia/Queensland	Australia/Brisbane
Australia/South	Australia/Adelaide
Australia/Sydney	Australia/Sydney
Australia/Tasmania	Australia/Hobart
Australia/Victoria	Australia/Sydney
Australia/West	Australia/Perth
Australia/Yancowinna	Australia/Broken_Hill
Brazil/Acre	America/Porto_Acre
Brazil/DeNoronha	America/Noronha
Brazil/East	America/Sao_Paulo
Brazil/West	America/Manaus
CET	Europe/Paris
CST	America/Chicago
CST6CDT	America/Chicago
Canada/Atlantic	America/Halifax
Canada/Central	America/Winnipeg
Canada/East-Saskatchewan	America/Regina
Canada/Eastern	America/Montreal
Canada/Mountain	America/Edmonton
Canada/Newfoundland	America/St_Johns
Canada/Pacific	America/Vancouver
Canada/Saskatchewan	America/Regina
Canada/Yukon	America/Los_Angeles
Chile/Continental	America/Santiago
Chile/EasterIsland	Pacific/Easter
Cuba	America/Havana
EET	America/Indianapolis
EST5EDT	America/New_York
Egypt	Africa/Cairo
Eire	Europe/Dublin
Etc/GMT	GMT
Etc/GMT0	GMT
Etc/Greenwich	GMT

System value	Java string object
Etc/UCT	UTC
Etc/UTC	UTC
Etc/Universal	UTC
Etc/Zulu	UTC
Europe/Amsterdam	Europe/Amsterdam
Europe/Andorra	Europe/Andorra
Europe/Athens	Europe/Athens
Europe/Belfast	Europe/London
Europe/Belgrade	Europe/Belgrade
Europe/Berlin	Europe/Berlin
Europe/Bratislava	Europe/Prague
Europe/Brussels	Europe/Brussels
Europe/Bucharest	Europe/Bucharest
Europe/Budapest	Europe/Budapest
Europe/Chisinau	Europe/Chisinau
Europe/Copenhagen	Europe/Copenhagen
Europe/Dublin	Europe/Dublin
Europe/Gibraltar	Europe/Gibraltar
Europe/Helsinki	Europe/Helsinki
Europe/Istanbul	Europe/Istanbul
Europe/Kaliningrad	Europe/Kaliningrad
Europe/Kiev	Europe/Kiev
Europe/Lisbon	Europe/Lisbon
Europe/Ljubljana	Europe/Belgrade
Europe/London	Europe/London
Europe/Luxembourg	Europe/Luxembourg
Europe/Madrid	Europe/Madrid
Europe/Malta	Europe/Malta
Europe/Minsk	Europe/Minsk
Europe/Monaco	Europe/Monaco
Europe/Moscow	Europe/Moscow
Europe/Oslo	Europe/Oslo
Europe/Paris	Europe/Paris
Europe/Prague	Europe/Prague
Europe/Riga	Europe/Riga
Europe/Rome	Europe/Rome
Europe/Samara	Europe/Samara
Europe/San_Marino	Europe/Rome
Europe/Sarajevo	Europe/Belgrade
Europe/Simferopol	Europe/Simferopol
Europe/Skopje	Europe/Belgrade

System value	Java string object
Europe/Sofia	Europe/Sofia
Europe/Stockholm	Europe/Stockholm
Europe/Tallinn	Europe/Tallinn
Europe/Tirane	Europe/Tirane
Europe/Vaduz	Europe/Vaduz
Europe/Vatican	Europe/Rome
Europe/Vienna	Europe/Vienna
Europe/Vilnius	Europe/Vilnius
Europe/Warsaw	Europe/Warsaw
Europe/Zagreb	Europe/Belgrade
Europe/Zurich	Europe/Zurich
Factory	GMT
GB	Europe/London
GB-Eire	Europe/London
GMT	GMT
GMT0	GMT
Greenwich	GMT
HST	Pacific/Honolulu
Hongkong	Asia/Hong_Kong
Iceland	Atlantic/Reykjavik
Indian/Antananarivo	Indian/Antananarivo
Indian/Chagos	Indian/Chagos
Indian/Christmas	Indian/Christmas
Indian/Cocos	Indian/Cocos
Indian/Comoro	Indian/Comoro
Indian/Kerguelen	Indian/Kerguelen
Indian/Mahe	Indian/Mahe
Indian/Maldives	Indian/Maldives
Indian/Mauritius	Indian/Mauritius
Indian/Mayotte	Indian/Mayotte
Indian/Reunion	Indian/Reunion
Iran	Asia/Tehran
Israel	Asia/Jerusalem
Jamaica	America/Jamaica
Japan	Asia/Tokyo
Libya	Africa/Tripoli
MET	Europe/Paris
MST	America/Phoenix
MST7MDT	America/Denver
Mexico/BajaNorte	America/Tijuana
Mexico/BajaSur	America/Mazatlan

System value	Java string object
Mexico/General	America/Mexico_City
NZ	Pacific/Auckland
NZ-CHAT	Pacific/Chatham
Navajo	America/Denver
PRC	Asia/Shanghai
PST	America/Los_Angeles
PST8PDT	America/Los_Angeles
Pacific/Apia	Pacific/Apia
Pacific/Auckland	Pacific/Auckland
Pacific/Chatham	Pacific/Chatham
Pacific/Easter	Pacific/Easter
Pacific/Efate	Pacific/Efate
Pacific/Enderbury	Pacific/Enderbury
Pacific/Fakaofu	Pacific/Fakaofu
Pacific/Fiji	Pacific/Fiji
Pacific/Funafuti	Pacific/Funafuti
Pacific/Galapagos	Pacific/Galapagos
Pacific/Gambier	Pacific/Gambier
Pacific/Guadalcanal	Pacific/Guadalcanal
Pacific/Guam	Pacific/Guam
Pacific/Honolulu	Pacific/Honolulu
Pacific/Kiritimati	Pacific/Kiritimati
Pacific/Kosrae	Pacific/Kosrae
Pacific/Majuro	Pacific/Majuro
Pacific/Marquesas	Pacific/Marquesas
Pacific/Nauru	Pacific/Nauru
Pacific/Niue	Pacific/Niue
Pacific/Norfolk	Pacific/Norfolk
Pacific/Noumea	Pacific/Noumea
Pacific/Pago_Pago	Pacific/Pago_Pago
Pacific/Palau	Pacific/Palau
Pacific/Pitcairn	Pacific/Pitcairn
Pacific/Ponape	Pacific/Ponape
Pacific/Port_Moresby	Pacific/Port_Moresby
Pacific/Rarotonga	Pacific/Rarotonga
Pacific/Saipan	Pacific/Saipan
Pacific/Samoa	Pacific/Pago_Pago
Pacific/Tahiti	Pacific/Tahiti
Pacific/Tarawa	Pacific/Tarawa
Pacific/Tongatapu	Pacific/Tongatapu
Pacific/Truk	Pacific/Truk

System value	Java string object
Pacific/Wake	Pacific/Wake
Pacific/Wallis	Pacific/Wallis
Poland	Europe/Warsaw
Portugal	Europe/Lisbon
ROC	Asia/Taipei
ROK	Asia/Seoul
Singapore	Asia/Singapore
SystemV/AST4ADT	America/Thule
SystemV/CST6CDT	America/Chicago
SystemV/EST5EDT	America/New_York
SystemV/MST7MDT	America/Denver
SystemV/PST8PDT	America/Los_Angeles
SystemV/YST9YDT	America/Anchorage
Turkey	Europe/Istanbul
UCT	UTC
US/Alaska	America/Anchorage
US/Aleutian	America/Adak
US/Arizona	America/Phoenix
US/Central	America/Chicago
US/East-Indiana	America/Indianapolis
US/Eastern	America/New_York
US/Hawaii	Pacific/Honolulu
US/Michigan	America/New_York
US/Mountain	America/Denver
US/Pacific	America/Los_Angeles
US/Pacific-New	America/Los_Angeles
US/Samoa	Pacific/Pago_Pago
UTC	UTC
Universal	UTC
W-SU	Europe/Moscow
WET	Europe/London
Zulu	UTC

Java locales

A locale is a geographic or political region of the world that shares the same language and customs. In Java^(TM), the Locale class represents a locale.

Supported Java locales

The IBM Developer Kit for Java supports these locales. The iSeries jobs and CNTRYID determine the default locale. See Java system properties for more details.

Locale Name in JDK 1.1.6	ISO Locale Name	Language/Country or region
ar	ar_EG	Arabic/Egypt
be	be_BY	Belorussian/Belarus
bg	bg_BG	Bulgarian/Bulgaria
ca	ca_ES	Catalan/Spain
cs	cs_CZ	Czech/Czech Republic
da	da_DK	Danish/Denmark
de	de_DE	German/Germany
de_AT	de_AT	German/Austria
de_CH	de_CH	German/Switzerland
el	el_GR	Greek/Greece
en	en_US	English/USA
en_AU	en_AU	English/Australia
en_CA	en_CA	English/Canada
en_GB	en_GB	English/Great Britain
en_IE	en_IE	English/Ireland
en_NZ	en_NZ	English/New Zealand
en_ZA	en_ZA	English/South Africa
es	es_ES	Spanish/Spain
es_AR	es_AR	Spanish/Argentina
es_BO	es_BO	Spanish/Bolivia
es_CL	es_CL	Spanish/Chile
es_CR	es_CR	Spanish/Costa Rica
es_DO	es_DO	Spanish/Dominican Republic
es_EC	es_EC	Spanish/Ecuador
es_GT	es_GT	Spanish/Guatemala
es_HN	es_HN	Spanish/Honduras
es_MX	es_MX	Spanish/Mexico
es_NI	es_NI	Spanish/Nicaragua
es_PA	es_PA	Spanish/Panama
es_PE	es_PE	Spanish/Peru
es_PR	es_PR	Spanish/Puerto Rico
es_PY	es_PY	Spanish/Paraguay
es_SV	es_SV	Spanish/El Salvador
es_UY	es_UY	Spanish/Uruguay
es_VE	es_VE	Spanish/Venezuela
et	et_EE	Estonian/Estonia
fi	fi_FI	Finnish/Finland
fr	fr_FR	French/France
fr_BE	fr_BE	French/Belgium
fr_CA	fr_CA	French/Canada
fr_CH	fr_CH	French/Switzerland

Locale Name in JDK 1.1.6	ISO Locale Name	Language/Country or region
hr	hr_HR	Croatian/Croatia
hu	hu_HU	Hungarian/Hungary
is	is_IS	Icelandic/Iceland
it	it_IT	Italian/Italy
it_CH	it_CH	Italian/Switzerland
iw	iw_IL	Hebrew/Israel
ja	ja_JP	Japanese/Japan
ko	ko_KR	Korean/Korea
lt	lt_LT	Lithuanian/Lithuania
lv	lv_LV	Latvian/Latvia
mk	mk_MK	Macedonian/Macedonia
nl	nl_NL	Dutch/Netherlands
nl_BE	nl_BE	Dutch/Belgium
no	no_NO_B	Norwegian/Norway
no_NO_NY	no_NO_NY	Norwegian/Norway, NY
pl	pl_PL	Polish/Poland
pt	pt_PT	Portuguese/Portugal
ro	ro_RO	Romanian/Romania
ru	ru_RU	Russian/Russia
sh	sh_SP	Serbo-Croatian/Serbia
sk	sk_SK	Slovak/Slovakia
sl	sl_SI	Slovene/Slovenian
sq	sq_AL	Albanian/Albania
sr	sr_SP	Serbian/Serbia
sv	sv_SE	Swedish/Sweden
tr	tr_TR	Turkish/Turkey
uk	uk_UA	Ukrainian/Ukraine
zh	zh_CN	Simplified Chinese
zh_TW	zh_TW	Traditional Chinese

Example: Internationalization of dates using the `java.util.DateFormat` class: This example shows how you can use locales to format dates.

Example 1: Demonstrates use of `java.util.DateFormat` class for internationalization of dates

Note: Read the Code example disclaimer for important legal information.

```
//*****
// File: DateExample.java
//*****

import java.text.*;
import java.util.*;
import java.util.Date;

public class DateExample {
```

```

public static void main(String args[]) {

    // Get the Date
    Date now = new Date();

    // Get date formatters for default, German, and French locales
    DateFormat theDate = DateFormat.getDateInstance(DateFormat.LONG);
    DateFormat germanDate = DateFormat.getDateInstance(DateFormat.LONG, Locale.GERMANY);
    DateFormat frenchDate = DateFormat.getDateInstance(DateFormat.LONG, Locale.FRANCE);

    // Format and print the dates
    System.out.println("Date in the default locale: " + theDate.format(now));
    System.out.println("Date in the German locale : " + germanDate.format(now));
    System.out.println("Date in the French locale : " + frenchDate.format(now));
}
}

```

For more information, see [Create an internationalized JavaTM program](#).

Example: Internationalization of numeric display using the `java.util.NumberFormat` class: This example shows how you can use locales to format numbers.

Example 1: Demonstrates use of `java.util.NumberFormat` class for internationalization of numeric output

Note: Read the Code example disclaimer for important legal information.

```

//*****
// File: NumberExample.java
//*****

import java.lang.*;
import java.text.*;
import java.util.*;

public class NumberExample {

    public static void main(String args[]) throws NumberFormatException {

        // The number to format
        double number = 12345.678;

        // Get formatters for default, Spanish, and Japanese locales
        NumberFormat defaultFormat = NumberFormat.getInstance();
        NumberFormat spanishFormat = NumberFormat.getInstance(new
Locale("es", "ES"));
        NumberFormat japaneseFormat = NumberFormat.getInstance(Locale.JAPAN);

        // Print out number in the default, Spanish, and Japanese formats
        // (Note: NumberFormat is not necessary for the default format)
        System.out.println("The number formatted for the default locale; " +
            defaultFormat.format(number));
        System.out.println("The number formatted for the Spanish locale; " +
            spanishFormat.format(number));
        System.out.println("The number formatted for the Japanese locale; " +
            japaneseFormat.format(number));
    }
}

```

For more information, see [Create an internationalized JavaTM program](#).

Example: Internationalization of locale-specific data using the `java.util.ResourceBundle` class: This example shows how you can use locales with resource bundles to internationalize program strings.

These property files are required for the `ResourceBundleExample` program to work as intended:

Contents of RBExample.properties

Hello.text=Hello

Contents of RBExample_de.properties

Hello.text=Guten Tag

Contents of RBExample_fr_FR.properties

Hello.text=Bonjour

Example 1: Demonstrates use of java.util.ResourceBundle class for internationalization of locale-specific data

Note: Read the Code example disclaimer for important legal information.

```
/**
 * *****
 * // File: ResourceBundleExample.java
 * *****
 */

import java.util.*;

public class ResourceBundleExample {
    public static void main(String args[]) throws MissingResourceException {

        String resourceName = "RBExample";
        ResourceBundle rb;

        // Default locale
        rb = ResourceBundle.getBundle(resourceName);
        System.out.println("Default : " + rb.getString("Hello" + ".text"));

        // Request a resource bundle with explicitly specified locale
        rb = ResourceBundle.getBundle(resourceName, Locale.GERMANY);
        System.out.println("German : " + rb.getString("Hello" + ".text"));

        // No property file for China in this example... use default
        rb = ResourceBundle.getBundle(resourceName, Locale.CHINA);
        System.out.println("Chinese : " + rb.getString("Hello" + ".text"));

        // Here is another way to do it...
        Locale.setDefault(Locale.FRANCE);
        rb = ResourceBundle.getBundle(resourceName);
        System.out.println("French : " + rb.getString("Hello" + ".text"));

        // No property file for China in this example... use default, which is now fr_FR.
        rb = ResourceBundle.getBundle(resourceName, Locale.CHINA);
        System.out.println("Chinese : " + rb.getString("Hello" + ".text"));
    }
}
```

For more information, see [Create an internationalized Java[™] program](#).

Java character encodings

Internally, the Java[™] virtual machine (JVM) always operates with data in Unicode. However, all data transferred into or out of the JVM is in a format matching the file.encoding property. Data read into the JVM is converted from file.encoding to Unicode and data sent out of the JVM is converted from Unicode to file.encoding.

Data files for Java programs are stored in the integrated file system. Files in the integrated file system are tagged with a coded character set identifier (CCSID) that identifies the character encoding of the data contained in the file. See the File.encoding values and iSeries CCSID table for description of how file.encoding is correlated to CCSID on the iSeries server.

When data is read by a Java program, it is expected to be in the character encoding matching file.encoding. When data is written to a file by a Java program, it is written in a character encoding

matching file.encoding. This also applies to Java source code files (.java files) processed by the javac command and to data sent and received through Transmission Control Protocol/Internet Protocol (TCP/IP) sockets using the .net package.

Data read from or written to System.in, System.out, and System.err are handled differently than data read from or written to other sources when they are assigned to stdin, stdout, and stderr. Since stdin, stdout, and stderr are normally attached to EBCDIC devices on the iSeries server, a conversion is performed by the JVM on the data to convert from the normal character encoding of file.encoding to a CCSID matching the iSeries job CCSID. When System.in, System.out, or System.err are redirected to a file or socket and are not directed to stdin, stdout, or stderr, this additional data conversion is not performed and the data remains in a character encoding matching file.encoding.

When data must be read into or written from a Java program in a character encoding other than file.encoding, the program can use the Java IO classes java.io.InputStreamReader, java.io.FileReader, java.io.OutputStreamReader, and java.io.FileWriter. These Java classes allow specifying a file.encoding value that takes precedence over the default file.encoding property currently in use by the JVM.

Data to or from the DB2/400 database, through the JDBC APIs, converts to or from the CCSID of the iSeries database.

Data that is transferred to or from other programs through Java Native Interface does not get converted.

For more information about internationalization, see OS/400 globalization.

You can also see Internationalization by Sun Microsystems, Inc. for more information.

File.encoding values and iSeries CCSID: This table shows the relation between possible file.encoding values and the closest matching iSeries coded character set identifier (CCSID).

file.encoding	CCSID	Description
Big5	950	8-bit ASCII T-Chinese BIG-5
CNS11643	964	Chinese National Character Set for traditional Chinese
Cp037	037	IBM EBCDIC US, Canada, Netherlands,...
Cp273	273	IBM EBCDIC Germany, Austria
Cp277	277	IBM EBCDIC Denmark, Norway
Cp278	278	IBM EBCDIC Finland, Sweden
Cp280	280	IBM EBCDIC Italy
Cp284	284	IBM EBCDIC Spanish, Latin America
Cp285	285	IBM EBCDIC UK
Cp297	297	IBM EBCDIC France
Cp420	420	IBM EBCDIC Arabic
Cp424	424	IBM EBCDIC Hebrew
Cp437	437	8-bit ASCII US PC
Cp500	500	IBM EBCDIC International
Cp737	737	8-bit ASCII Greek MS-DOS
Cp775	775	8-bit ASCII Baltic MS-DOS
Cp838	838	IBM EBCDIC Thailand
Cp850	850	8-bit ASCII Latin-1 Multinational

file.encoding	CCSID	Description
Cp852	852	8-bit ASCII Latin-2
Cp855	855	8-bit ASCII Cyrillic
Cp856	856	8-bit ASCII Hebrew
Cp857	857	8-bit ASCII Latin-5
Cp860	860	8-bit ASCII Portugal
Cp861	861	8-bit ASCII Iceland
Cp862	862	8-bit ASCII Hebrew
Cp863	863	8-bit ASCII Canada
Cp864	864	8-bit ASCII Arabic
Cp865	865	8-bit ASCII Denmark, Norway
Cp866	866	8-bit ASCII Cyrillic
Cp868	868	8-bit ASCII Urdu
Cp869	869	8-bit ASCII Greek
Cp870	870	IBM EBCDIC Latin-2
Cp871	871	IBM EBCDIC Iceland
Cp874	874	8-bit ASCII Thailand
Cp875	875	IBM EBCDIC Greek
Cp918	918	IBM EBCDIC Urdu
Cp921	921	8-bit ASCII Baltic
Cp922	922	8-bit ASCII Estonia
Cp930	930	IBM EBCDIC Japanese Extended Katakana
Cp933	933	IBM EBCDIC Korean
Cp935	935	IBM EBCDIC Simplified Chinese
Cp937	937	IBM EBCDIC Traditional Chinese
Cp939	939	IBM EBCDIC Japanese Extended Latin
Cp942	942	8-bit ASCII Japanese
Cp943	943	Japanese PC data mixed for open env
Cp943C	943	Japanese PC data mixed for open env
Cp948	948	8-bit ASCII IBM Traditional Chinese
Cp949	949	8-bit ASCII Korean KSC5601
Cp950	950	8-bit ASCII T-Chinese BIG-5
Cp964	964	EUC Traditional Chinese
Cp970	970	EUC Korean
Cp1006	1006	ISO 8-bit Urdu
Cp1025	1025	IBM EBCDIC Cyrillic
Cp1026	1026	IBM EBCDIC Turkey
Cp1046	1046	8-bit ASCII Arabic
Cp1097	1097	IBM EBCDIC Farsi

file.encoding	CCSID	Description
Cp1098	1098	8-bit ASCII Farsi
Cp1112	1112	IBM EBCDIC Baltic
Cp1122	1122	IBM EBCDIC Estonia
Cp1123	1123	IBM EBCDIC Ukraine
Cp1124	1124	ISO 8-bit Ukraine
Cp1250	1250	MS-Win Latin-2
Cp1251	1251	MS-Win Cyrillic
Cp1252	1252	MS-Win Latin-1
Cp1253	1253	MS-Win Greek
Cp1254	1254	MS-Win Turkish
Cp1255	1255	MS-Win Hebrew
Cp1256	1256	MS-Win Arabic
Cp1257	1257	MS-Win Baltic
Cp1258	1251	MS-Win Russian
Cp1381	1381	8-bit ASCII S-Chinese GB
Cp1383	1383	EUC Simplified Chinese
Cp33722	33722	EUC Japanese
EUC_CN	1383	EUC for Simplified Chinese
EUC_JP	33722	EUC for Japanese
EUC_KR	970	EUC for Korean
EUC_TW	964	EUC for Traditional Chinese
GB2312	1381	8-bit ASCII S-Chinese GB
GBK	1386	New simplified Chinese 8-bit ASCII 9
ISO2022CN_CNS	None Available	7-bit ASCII for Traditional Chinese
ISO2022CN_GB	None Available	7-bit ASCII for Simplified Chinese
ISO2022JP	5054	7-bit ASCII for Japanese
ISO2022KR	25546	7-bit ASCII for Korean
ISO8859_1	819	ISO 8859-1 ISO Latin-1
ISO8859_2	912	ISO 8859-2 ISO Latin-2
ISO8859_3	913	ISO 8859-3 ISO Latin-3
ISO8859_4	914	ISO 8859-4 ISO Latin-4
ISO8859_5	915	ISO 8859-5 ISO Latin-5
ISO8859_6	1089	ISO 8859-6 ISO Latin-6 (Arabic)
ISO8859_7	813	ISO 8859-7 ISO Latin-7 (Greek/Latin)
ISO8859_8	916	ISO 8859-8 ISO Latin-8 (Hebrew)
ISO8859_9	920	ISO 8859-9 ISO Latin-9 (ECMA-128, Turkey)
JIS0201	897	Japanese industry standard X0201
JIS0208	952	Japanese industry standard X0208
JIS0212	953	Japanese industry standard X0212

file.encoding	CCSID	Description
Johab	None Available	Korean composition Hangul encoding (full)
K018_R		Cyrillic
KSC5601	949	8-bit ASCII Korean
MS874	874	MS-Win Thailand
SJIS	932	8-bit ASCII Japanese
TIS620	874	Thai industry standard 620
UTF8	1208	UTF-8 (IBM CCSID 1208, which is not yet available on the iSeries server)
Unicode	13488	UNICODE, UCS-2
UnicodeBig	13488	Same as Unicode
UnicodeBigUnmarked		Unicode with no byte-order mark
UnicodeLittle		Unicode with little-endian byte order
UnicodeLittleUnmarked		UnicodeLittle with no byte-order mark






For default values, see Default file.encoding values.

Default file.encoding values: This table shows how the file.encoding value is set based on the iSeries coded character set identifier (CCSID) when the Java^(TM) virtual machine starts.

iSeries CCSID	Default file.encoding	Description
37	ISO8859_1	English for USA, Canada, New Zealand, and Australia; Portuguese for Portugal and Brazil; and Dutch for Netherlands
256	ISO8859_1	International #1
273	ISO8859_1	German/Germany, German/Austria
277	ISO8859_1	Danish/Denmark, Norwegian/Norway, Norwegian/Norway, NY
278	ISO8859_1	Finnish/Finland
280	ISO8859_1	Italian/Italy
284	ISO8859_1	Catalan/Spain, Spanish/Spain
285	ISO8859_1	English/Great Britain, English/Ireland
290	Cp943C	SBCS portion of Japanese EBCDIC mixed (CCSID 5026)
297	ISO8859_1	French/France
420	Cp1046	Arabic/Egypt
423	ISO8859_7	Greece
424	ISO8859_8	Hebrew/Israel
500	ISO8859_1	German/Switzerland, French/Belgium, French/Canada, French/Switzerland
833	Cp970	SBCS portion of Korean EBCDIC mixed (CCSID 933)

iSeries CCSID	Default file.encoding	Description
836	Cp1383	SBCS portion of S-Chinese EBCDIC mixed (CCSID 935).
838	TIS620	Thai
870	ISO8859_2	Czech/Czech Republic, Croatian/Croatia, Hungarian/Hungary, Polish/Poland
871	ISO8859_1	Icelandic/Iceland
875	ISO8859_7	Greek/Greece
880	ISO8859_5	Bulgaria (ISO 8859_5)
905	ISO8859_9	Turkey extended
918	Cp868	Urdu
930	Cp943C	Japanese EBCDIC mixed (similar to CCSID 5026)
933	Cp970	Korean/Korea
935	Cp1383	Simplified Chinese
937	Cp950	Traditional Chinese
939	Cp943C	Japanese EBCDIC Mixed (similar to CCSID 5035)
1025	ISO8859_5	Belorussian/Belarus, Bulgarian/Bulgaria, Macedonian/Macedonia, Russian/Russia
1026	ISO8859_9	Turkish/Turkey
1027	Cp943C	SBCS portion of Japanese EBCDIC mixed (CCSID 5035)
1097	Cp1098	Farsi
1112	Cp921	Lithuanian/Lithuania, Latvian/Latvia, Baltic
1388	GBK	Simplified Chinese EBCDIC mixed (GBK is included)
5026	Cp943C	Japanese EBCDIC mixed CCSID (Extended Katakana)
5035	Cp943C	Japanese EBCDIC mixed CCSID (Extended Latin)
8612	Cp1046	Arabic (base shapes only) (or ASCII 420 and 8859_6)
9030	Cp874	Thai (host extended SBCS)
13124	GBK	SBCS portion of Simplified Chinese EBCDIC mixed (GBK is included)
28709	Cp948	SBCS portion of Traditional Chinese EBCDIC mixed (CCSID 937)

Release-to-release compatibility

Java^(TM) class files are upward compatible (JDK 1.1.x -> 1.2.x -> 1.3.x  -> 1.4.x ) as long as they do not make use of a few features that Sun has dropped or changed support for (See Sun documentation). Class files are also downward compatible ( 1.4.x -> 1.3.x -> 1.2.x -> 1.1.x ) as long as the programs only make use of Java features that are available in the earlier JDK level. This holds true on the iSeries server as well, since it is a compliant implementation. See The Source for Java Technology java.sun.com  for information on release to release availability.



When Java programs on an iSeries server are optimized using Create Java Program (CRTJVAPGM) command, a JVAPGM is attached to the class file. The internal structure of these JVAPGMs changed on V4R4. This means that JVAPGMs created before V4R4 are not be valid on V4R4 and later releases. They must to be recreated. If no action is taken, the system automatically creates a JVAPGM at the same optimization level as before. It is, however, recommended that you manually do a CRTJVAPGM, especially with JAR or ZIP files. This produces the best optimization with the smallest program size.

For best performance at optimization level 40, it is recommended to do CRTJVAPGM on each OS/400 release or JDK version change. This is especially true if the JDKVER facility is used on CRTJVAPGM, as this results in methods from the Sun JDK being inlined into your JVAPGM. This can be a great advantage to performance. If, however, changes are made in the JDK on subsequent releases that invalidate those inlines, the programs may actually run slower than at lower optimizations. This is because special case code must be run to get proper operation.


See Java runtime performance for more detailed performance information.

Database access with the IBM Developer Kit for Java

With the IBM Developer Kit for Java^(TM), your Java programs can access your database files in three ways:

- JDBC driver explains how the IBM Developer Kit for Java JDBC driver allows Java programs to access database files.
- SQLJ support explains how the IBM Developer Kit for Java allows you to use SQL statements that are embedded in your Java application.
-  Java SQL routines explains how you can use Java stored procedures and Java user-defined functions to access Java programs. 

Access your iSeries database with the IBM Developer Kit for Java JDBC driver

The IBM Developer Kit for Java^(TM) JDBC driver, also known as the "native" driver, provides programmatic access to iSeries database files. Using the Java Database Connectivity (JDBC) API, applications written in the Java language can access JDBC database functions with embedded Structured Query Language (SQL), run SQL statements, retrieve results, and propagate changes back to the database.  The JDBC API can also be used to interact with multiple data sources in a distributed, heterogeneous environment.

The SQL99 Command Language Interface (CLI), on which the JDBC API is based, is the basis for ODBC. JDBC provides a natural and easy-to-use mapping from the Java programming language to the abstractions and concepts defined in the SQL standard.

To use the JDBC driver, see the following:

Get started with JDBC

You can follow the tutorial of writing a JDBC program and running it on your iSeries server.

Connections

An application program can have multiple connections at one time. You can represent a connection to a data source in JDBC by using a Connection object. It is through Connection objects that Statement objects are created for processing SQL statements against the database.

DatabaseMetaData

The DatabaseMetaData interface is used by application servers and tools to determine how to interact with a given data source. Applications may also use DatabaseMetaData methods to obtain information about a specific data source.

Exceptions

The Java language uses exceptions to provide error-handling capabilities for its programs. An exception is an event that occurs when you run your program that disrupts the normal flow of instructions.

Transactions

A transaction is a logical unit of work. Transactions are used to provide data integrity, correct application semantics, and a consistent view of data during concurrent access. All JDBC-compliant drivers must support transactions.

Statement types

The Statement interface and its PreparedStatement and CallableStatement subclasses are used to process SQL commands against the database. SQL statements cause the generation of ResultSet objects.

ResultSets

The ResultSet interface provides access to the results generated by running queries. Data of a ResultSet can be thought of as a table with a specific number of columns and a specific number of rows. By default, the table rows are retrieved in sequence. Within a row, column values can be accessed in any order.

JDBC object pooling

Since many objects used in JDBC are expensive to create such as Connection, Statement, and ResultSet objects, significant performance benefits can be achieved by using JDBC object pooling. With object pooling, you can reuse these objects instead of creating them every time you need them.

Batch updates

Batch update support allows many updates to the database to be passed as a single transaction between the user program and the database. Batch updates can significantly improve performance when many updates must be performed at once.

Advanced data types

There are several new data types called SQL3 data types that are provided in the iSeries database. The SQL3 data types give you a tremendous amount of flexibility. They are ideal for storing serialized Java objects, Extensible Markup Language (XML) documents, and multimedia data such as songs, product pictures, employee photographs, and movie clips. The SQL3 data types include the following:

- Distinct types
- Large objects such as Binary Large Objects, Character Large Objects, and Double Byte Character Large Objects
- Datalinks

RowSets

The RowSet specification is designed to be more of a framework than an actual implementation. The RowSet interfaces define a set of core functionality that all RowSets have.



Distributed transactions

The Java Transaction API (JTA) has support for complex transactions. It also provides support for decoupling transactions from Connection objects. JTA and JDBC work together to decouple transactions from Connection objects and allows you to have a single connection work on multiple transactions concurrently. Conversely, it allows you to have multiple connections work on a single transaction.


Performance tips

You can obtain the best possible performance from your JDBC applications with these performance tips.

For more information about JDBC, see the JDBC  documentation by Sun Microsystems, Inc.

For more information about iSeries Native JDBC driver, see IBM Developer Kit for Java JDBC Web Page  . 

Get started with JDBC

 The Java^(TM) Database Connectivity (JDBC) driver shipped with the Developer Kit for Java is called the Developer Kit for Java JDBC driver. This driver is also commonly known as the native JDBC driver.

To select which JDBC driver suits your needs, consider the following suggestions:

- Programs running directly on a server where the database resides should use the native JDBC driver for performance. This includes most servlet and JavaServer Pages (JSP) solutions, and applications written to run locally on an iSeries server.
- Programs that must connect to a remote iSeries server use the Toolbox JDBC driver. The Toolbox is a robust implementation of JDBC and is provided as part of the Toolbox for Java. Being pure Java, the Toolbox JDBC driver is trivial to set up for clients and requires little server setup.
- Programs that run on an iSeries server and need to connect to a remote, non-iSeries database use the native JDBC driver and set up a Distributed Relational Database Architecture (DRDA) connection to that remote server.

To get started with JDBC, see the following:

Types of JDBC drivers

This topic defines the JDBC driver types. Driver types are defined to categorize the technology used to connect to the database.

Requirements


This topic indicates the requirements you need to access the following:

- Core JDBC
- JDBC 2.0 optional package
- Java Transaction API (JTA)

JDBC tutorial

This is an important first step towards writing a JDBC program and having it run on an iSeries server with the native JDBC driver.



Types of JDBC drivers:  This topic defines the Java^(TM) Database Connectivity (JDBC) driver types. Driver types are used to categorize the technology used to connect to the database. A JDBC driver vendor uses these types to describe how their product operates. Some JDBC driver types are better suited for some applications than others.

Type 1: Type 1 drivers are "bridge" drivers. They use another technology such as Open Database Connectivity (ODBC) to communicate with a database. This is an advantage because ODBC drivers exist for many Relational Database Management System (RDBMS) platforms. The Java Native Interface (JNI) is used to call ODBC functions from the JDBC driver.

A Type 1 driver needs to have the bridge driver installed and configured before JDBC can be used with it. This can be a serious drawback for a production application. Type 1 drivers cannot be used in an applet since applets cannot load native code.

Type 2: Type 2 drivers use a native API to communicate with a database system. Java native methods are used to invoke the API functions that perform database operations. Type 2 drivers are generally faster than Type 1 drivers.

Type 2 drivers need native binary code installed and configured to work. A Type 2 driver also uses the JNI. You cannot use a Type 2 driver in an applet since applets cannot load native code. A Type 2 JDBC driver may require some Database Management System (DBMS) networking software to be installed.

The Developer Kit for Java JDBC driver is a Type 2 JDBC driver.

Type 3: These drivers use a networking protocol and middleware to communicate with a server. The server then translates the protocol to DBMS function calls specific to DBMS.

Type 3 JDBC drivers are the most flexible JDBC solution because they do not require any native binary code on the client. A Type 3 driver does not need any client installation.

Type 4: A Type 4 driver uses Java to implement a DBMS vendor networking protocol. Since the protocols are usually proprietary, DBMS vendors are generally the only companies providing a Type 4 JDBC driver.

Type 4 drivers are all Java drivers. This means that there is no client installation or configuration. However, a Type 4 driver may not be suitable for some applications if the underlying protocol does not handle issues such as security and network connectivity well.

The Toolbox JDBC driver is a Type 4 JDBC driver, indicating that the API is a pure Java networking protocol driver. <<

JDBC requirements: >> Before you write and deploy your JDBC applications, you may need to install the following:

- Core JDBC
- JDBC 2.0 optional package
- Java Transaction API

Core JDBC: For core Java^(TM) Database Connectivity (JDBC) access to the local database, there are no requirements. All support is built in, pre-installed, and configured.

Notes:

- In the past, you needed to add a symbolic link for JDK 1.2 and higher. This requirement has been removed with V4R5 PTF SF65439 and V5R1 PTF SI00959.
- In the past, you had to ensure that there was at least one entry in the Relation Database Directory to connect to a system (typically *LOCAL for connecting to the local database). This requirement has also been lifted. If there is no entry for the local database, the entry is created using the name of the system when the local system is accessed. Relation Database Directory entries still have to be configured for any remote systems you want to connect to with the native JDBC driver.

JDBC 2.0 optional package: If you need to use the classes of the JDBC 2.0 optional package, you must include the jdbc2_0-stdext.jar file in your classpath. This Java ARchive (JAR) file contains all the standard

interfaces that you need to write your application to use the JDBC 2.0 optional package. To add the JAR file to your extensions classpath, create a symbolic link from the UserData extensions directory to the location of the JAR file. You only need to perform this once; the JDBC 2.0 optional package JAR file is always available to your applications at runtime. Use the following command to add the optional package to the extensions classpath:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/Java400/ext/jdbc2_0-stdext.jar')
NEWLNK('/QIBM/UserData/Java400/ext/jdbc2_0-stdext.jar')
```

Note: This requirement is only for JDKs 1.2 and 1.3. Since JDK 1.4 is the first release with JDBC 3.0 support, all of JDBC (that is, the core JDBC and the optional package) moves into the base JDK runtime JAR file that you program always finds.


Java Transaction API: If you need to use the Java Transaction API (JTA) in your application, you must include the jta-spec1_0_1.jar file in your classpath. This JAR file contains all the standard interfaces that you need to write your application to use JTA. To add the JAR file to you extensions classpath, create a symbolic link from the UserData extensions directory to the location of the JAR file. This is a one-time operation and once completed, the JTA JAR file is always available to your application at runtime. Use the following command to add JTA to the extensions classpath:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/Java400/ext/jta-spec1_0_1.jar')
NEWLNK('/QIBM/UserData/Java400/ext/jta-spec1_0_1.jar')
```

JDBC compliance: The native JDBC driver is compliant with all relevant JDBC specifications. The compliance level of the JDBC driver is not dependent on the OS/400 release, but on the JDK release you use. The native JDBC driver's compliance level for the various JDKs is listed as follows:

JDK release	JDBC driver's compliance level
JDK 1.1	This JDK is compliant with JDBC 1.0.
JDK 1.2	This JDK is compliant with JDBC 2.0 and supports JDBC 2.1 optional package.
JDK 1.3	This JDK is compliant with JDBC 2.0 and supports JDBC 2.1 optional package (there were no JDBC-related changes for JDK 1.3).
JDK 1.4	This JDK is compliant with JDBC 3.0, but the JDBC optional package no longer exists (support for it is now part of the core JDK).



JDBC tutorial:  The following is a tutorial on writing a Java^(TM) Database Connectivity (JDBC) program and having it run on the an iSeries server with the native JDBC driver. It is designed to show you the basic steps required for your program to run JDBC.

The example program creates a table and populates it with some data. The program processes a query to get that data out of the database and to display it on the screen.

Run the example program: To run the example program, perform the following steps:

1. Copy the program to your workstation.
 - a. Copy the example program and paste it into a file on your workstation.
 - b. Save the file with the same name as the public class provided and with the .java extension. In this case, you must name the file BasicJDBC.java on your local workstation.
2. Transfer the file from your workstation to your iSeries server. From a command prompt, enter the following commands:

```
ftp <iSeries server name>
<Enter your user ID>
<Enter your password>
cd /home/cujo
put BasicJDBC.java
quit
```

For these commands to work, you must have a directory in which to put the file. In the example, /home/cujo is the location, but you can use any location you want.

Note: It is possible that the FTP commands mentioned previously may be different for you based on how your server is set up, but they should be similar. It does not matter how you transfer the file to your iSeries server as long as you transfer it into the integrated file system. Tools such as VisualAge for Java can fully automate this process for you.

3. Make sure you set your classpath to the directory where you put the file in so that your Java commands find the file when you run them. From a CL command line, you can use WRKENVVAR to see what environment variables are set for your user profile.
 - If you see an environment variable named CLASSPATH, you must ensure that the location where you put the .java file in is in the string of directories listed there or add it if the location has not been specified.
 - If there is no CLASSPATH environment variable, you must add one. This can be accomplished with the following command:

```
ADDENVVAR ENVVAR(CLASSPATH)
VALUE('/home/cujo:/QIBM/ProdData/Java400/jdk13/lib/tools.jar')
```

Note: To compile Java code from the CL command, you must include the tools.jar file. This JAR file includes the javac command.

4. Compile the Java file into a class file.
Enter the following command from the CL command line:

```
java class(com.sun.tools.javac.Main) prop(BasicJDBC)
java BasicJDBC
```

You can also compile the Java file from QSH:

```
cd /home/cujo
javac BasicJDBC.java
```

QSH automatically ensures that the tools.jar file can be found. As a result, you do not have to add it to your classpath. The current directory is also in the classpath. By issuing the change directory (cd) command, the BasicJDBC.java file is also found.

Note: You can also compile the file on your workstation and use FTP to send the class file to your iSeries server in binary mode. This is an example of Java's ability to run on any platform. Run the program by using the following command from either the CL command line or from QSH:




```
java BasicJDBC
```

The output is as follows:

```
-----
| 1 | Frank Johnson |
| 2 | Neil Schwartz  |
| 3 | Ben Rodman     |
| 4 | Dan Gloore     |
```

```
-----  
There were 4 rows returned.  
Output is complete.  
Java program completed.
```

References: For more information on Java and JDBC, consult the following resources:

- Native JDBC driver external web site 
- Toolbox JDBC driver external web site 
- Sun's JDBC page 
- Java/JDBC forum for iSeries and iSeries users
- IBM JDBC e-mail address



Example: JDBC:  This is an example of how to use the BasicJDBC program.

Example: BasicJDBC

Note: Read the Code example disclaimer for important legal information.

```
/////////////////////////////////////////////////////////////////  
//  
// BasicJDBC example. This program uses the native JDBC driver for the  
// Developer Kit for Java to build a simple table and process a query  
// that displays the data in that table.  
//  
// Command syntax:  
//   BasicJDBC  
//  
/////////////////////////////////////////////////////////////////  
//  
// This source is an example of the IBM Developer for Java JDBC driver.  
// IBM grants you a nonexclusive license to use this as an example  
// from which you can generate similar function tailored to  
// your own specific needs.  
//  
// This sample code is provided by IBM for illustrative purposes  
// only. These examples have not been thoroughly tested under all  
// conditions. IBM, therefore, cannot guarantee or imply  
// reliability, serviceability, or function of these programs.  
//  
// All programs contained herein are provided to you "AS IS"  
// without any warranties of any kind. The implied warranties of  
// merchantability and fitness for a particular purpose are  
// expressly disclaimed.  
//  
// IBM Developer Kit for Java  
// (C) Copyright IBM Corp. 2001  
// All rights reserved.  
// US Government Users Restricted Rights -  
// Use, duplication, or disclosure restricted  
// by GSA ADP Schedule Contract with IBM Corp.  
//  
/////////////////////////////////////////////////////////////////  
  
// Include any Java classes that are to be used. In this application,  
// many classes from the java.sql package are used and the  
// java.util.Properties class is also used as part of obtaining  
// a connection to the database.  
import java.sql.*;  
import java.util.Properties;
```

```

// Create a public class to encapsulate the program.
public class BasicJDBC {

    // The connection is a private variable of the object.
    private Connection connection = null;

    // Any class that is to be an 'entry point' for running
    // a program must have a main method. The main method
    // is where processing begins when the program is called.
    public static void main(java.lang.String[] args) {

        // Create an object of type BasicJDBC. This
        // is fundamental to object-oriented programming. Once
        // an object is created, call various methods on
        // that object to accomplish work.
        // In this case, calling the constructor for the object
        // creates a database connection that the other
        // methods use to do work against the database.
        BasicJDBC test = new BasicJDBC();

        // Call the rebuildTable method. This method ensures that
        // the table used in this program exists and looks the
        // way it should. The return value is a boolean for
        // whether or not rebuilding the table completed
        // successfully. If it did no, display a message
        // and exit the program.
        if (!test.rebuildTable()) {
            System.out.println("Failure occurred while setting up " +
                " for running the test.");
            System.out.println("Test will not continue.");
            System.exit(0);
        }

        // The run query method is called next. This method
        // processes an SQL select statement against the table that
        // was created in the rebuildTable method. The output of
        // that query is output to standard out for you to view.
        test.runQuery();

        // Finally, the cleanup method is called. This method
        // ensures that the database connection that the object has
        // been hanging on to is closed.
        test.cleanup();
    }

    /**
    This is the constructor for the basic JDBC test. It creates a database
    connection that is stored in an instance variable to be used in later
    method calls.
    */
    public BasicJDBC() {

        // One way to create a database connection is to pass a URL
        // and a java Properties object to the DriverManager. The following
        // code constructs a Properties object that has your user ID and
        // password. These pieces of information are used for connecting
        // to the database.
        Properties properties = new Properties ();
        properties.put("user", "cujo");
        properties.put("user", "newtiger");

        // Use a try/catch block to catch all exceptions that can come out of the
        // following code.
        try {
            // The DriverManager must be aware that there is a JDBC driver available
            // to handle a user connection request. The following line causes the

```

```

// native JDBC driver to be loaded and registered with the DriverManager.
Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

// Create the database Connection object that this program uses in all
// the other method calls that are made. The following code specifies
// that a connection is to be established to the local database and that
// that connection should conform to the properties that were set up
// previously (that is, it should use the user ID and password specified).
connection = DriverManager.getConnection("jdbc:db2:*local", properties);

} catch (Exception e) {
// If any of the lines in the try/catch block fail, control transfers to
// the following line of code. A robust application tries to handle the
// problem or provide more details to you. In this program, the error
// message from the exception is displayed and the application allows
// the program to return.
System.out.println("Caught exception: " + e.getMessage());
}
}

/**
Ensures that the qgpl.basicjdbc table looks you want it to at the start of
the test.

@return boolean Returns true if the table was rebuild successfully;
returns false if any failure occurred.
**/
public boolean rebuildTable() {
// Wrap all the functionality in a try/catch block so an attempt is
// made to handle any errors that may happen within this method.
try {

// Statement objects are used to process SQL statements against the
// database. The Connection object is used to create a Statement
// object.
Statement s = connection.createStatement();

try {
// Build the test table from scratch. Process an update statement
// that attempts to delete the table if it currently exists.
s.executeUpdate("drop table qgpl.basicjdbc");
} catch (SQLException e) {
// Do not perform anything if an exception occurred. Assume
// that the problem is that the table that was dropped does not
// exist and that it can be created next.
}

// Use the statement object to create our table.
s.executeUpdate("create table qgpl.basicjdbc(id int, name char(15))");

// Use the statement object to populate our table with some data.
s.executeUpdate("insert into qgpl.basicjdbc values(1, 'Frank Johnson')");
s.executeUpdate("insert into qgpl.basicjdbc values(2, 'Neil Schwartz')");
s.executeUpdate("insert into qgpl.basicjdbc values(3, 'Ben Rodman')");
s.executeUpdate("insert into qgpl.basicjdbc values(4, 'Dan Gloore')");

// Close the SQL statement to tell the database that it is no longer
// needed.
s.close();

// If the entire method processed successfully, return true. At this point,
// the table has been created or refreshed correctly.
return true;

} catch (SQLException sqle) {
// If any of our SQL statements failed (other than the drop of the table

```

```

        // that was handled in the inner try/catch block), the error message is
        // displayed and false is returned to the caller, indicating that the table
        // may not be complete.
        System.out.println("Error in rebuildTable: " + sqle.getMessage());
        return false;
    }
}

/**
Runs a query against the demonstration table and the results are displayed to
standard out.
**/
public void runQuery() {
    // Wrap all the functionality in a try/catch block so an attempts is
    // made to handle any errors that might happen within this
    // method.
    try {
        // Create a Statement object.
        Statement s = connection.createStatement();

        // Use the statement object to run an SQL query. Queries return
        // ResultSet objects that are used to look at the data the query
        // provides.
        ResultSet rs = s.executeQuery("select * from qgpl.basicjdbc");

        // Display the top of our 'table' and initialize the counter for the
        // number of rows returned.
        System.out.println("-----");
        int i = 0;

        // The ResultSet next method is used to process the rows of a
        // ResultSet. The next method must be called once before the
        // first data is available for viewing. As long as next returns
        // true, there is another row of data that can be used.
        while (rs.next()) {

            // Obtain both columns in the table for each row and write a row to
            // our on-screen table with the data. Then, increment the count
            // of rows that have been processed.
            System.out.println("| " + rs.getInt(1) + " | " + rs.getString(2) + "|");
            i++;
        }

        // Place a border at the bottom on the table and display the number of rows
        // as output.
        System.out.println("-----");
        System.out.println("There were " + i + " rows returned.");
        System.out.println("Output is complete.");
    } catch (SQLException e) {
        // Display more information about any SQL exceptions that are
        // generated as output.
        System.out.println("SQLException exception: ");
        System.out.println("Message:....." + e.getMessage());
        System.out.println("SQLState:...." + e.getSQLState());
        System.out.println("Vendor Code:.." + e.getErrorCode());
        e.printStackTrace();
    }
}

/**
The following method ensures that any JDBC resources that are still
allocated are freed.
**/

```



```

public void cleanup() {
    try {
        if (connection != null)
            connection.close();
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}
}


```



Use JNDI for the examples: ➤ DataSources work hand-in-hand with the Java^(TM) Naming and Directory Interface (JNDI). JNDI is a Java abstraction layer for directory services just as Java Database Connectivity (JDBC) is an abstraction layer for databases. JNDI is used most often with the Lightweight Directory Access Protocol (LDAP), but it may also be used with the CORBA Object Services (COS), the Java Remote Method Invocation (RMI) registry, or the underlying file system. This varied use is accomplished by means of the various directory service providers that turn common JNDI requests into specific directory service requests.

The DataSource samples were designed using the JNDI file system service provider. If you want to run the examples provided, there must be a JNDI service provider in place.

Follow these directions to set up the environment for the file system service provider:

1. Download the file system JNDI support from Sun Microsystems JNDI site .
2. Click **Continue** to download JNDI 1.2.1. A license agreement displays.
3. Click **Accept** and then click **FS Context** to display the download options for JNDI context support.
4. Download fscontext.zip and unzip the file from your workstation.
5. Use FTP to transfer fscontext.jar and providerutil.jar to your system and put them in /QIBM/UserData/Java400/ext. This is the extensions directory and the JAR files that you place here are found automatically when you run your application (that is, you do not need them in your classpath).

Once you have support for a service provider for JNDI, you must set up the context information for your applications. This can be accomplished by putting the required information in a SystemDefault.properties file. There are several places on the system where you can specify default properties, but the best way is to create a text file called SystemDefault.properties in your home directory (that is, at /home/).

To create a file, use the following lines or add them to your existing file:

```

# Needed env settings for JNDI.
java.naming.factory.initial=com.sun.jndi.fscontext.RefFSContextFactory
java.naming.provider.url=file:/DataSources/jdbc

```

These lines specify that the file system service provider handles JNDI requests and that /DataSources/jdbc is the root for tasks that use JNDI. You can change this location, but the directory that you specify must exist. The location that you specify is where the example DataSources are bound and deployed. <<

Connections

➤ The Connection object represents a connection to a data source in Java^(TM) Database Connectivity (JDBC). It is through Connection objects that Statement objects are created for processing SQL statements against the database. An application program can have multiple connections at one time. These Connection objects can all connect to the same database or connect to different databases.

Obtaining a connection in JDBC can be accomplished in two ways:

- Through the DriverManager class.
- By using DataSources.

Using DataSources to obtain a connection is preferred because it enhances application portability and maintainability. It also allows an application to transparently use connection and statement pooling, and distributed transactions.

For details on obtaining connections, see the following sections:

DriverManager

The DriverManager is a static class that manages the set of available JDBC drivers for an application to use.

Connection properties

The table lists valid JDBC driver connection properties, their values, and their descriptions.

Use DataSources with UDBDataSource

You can deploy a DataSource with the UDBDataSource class by setting it up to have specific properties and then binding it into some directory service through the use of the Java Naming and Directory Interface (JNDI).

DataSource properties

The table lists valid DataSource properties, their values, and their descriptions.

Other DataSource implementations

There are other implementations of the DataSource interface provided with the native JDBC driver. They exist only to serve as a bridge until the UDBDataSource and its related functions are adopted.

Once a connection is obtained, it can be used to accomplish the following JDBC tasks:

- Create various types of Statement objects for interacting with the database.
- Control transactions against the database.
- Retrieve metadata about the database.



DriverManager: ➤ DriverManager is a static class provided by the Java^(TM) Developer Kit (JDK). It is responsible for managing the set of Java Database Connectivity (JDBC) drivers that are available for an application to use. Applications can use multiple JDBC drivers concurrently if necessary. How an application program specifies a JDBC driver is by using a Uniform Resource Locator (URL). By passing a URL for a specific JDBC driver to the DriverManager, the application informs the DriverManager about which type of JDBC connection should be returned to the application.

Before this can be done, the DriverManager must be made aware of the available JDBC drivers so it can hand out connections. By making a call to the Class.forName method, it loads a class into the running Java virtual machine (JVM) based on its string name that is passed into the method. The following is an example of the class.forName method being used to load the native JDBC driver:

Example: Load the native JDBC driver

Note: Read the Code example disclaimer for important legal information.

```
// Load the native JDBC driver into the DriverManager to make it
// available for getConnection requests.
```

```
Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
```

JDBC drivers are designed to tell the DriverManager about themselves automatically when their driver implementation class loads. Once the line of code previously mentioned has been processed, the native JDBC driver is available for the DriverManager with which to work. The following line of code requests a Connection object using the native JDBC URL:

Example: Request a Connection object

Note: Read the Code example disclaimer for important legal information.

```
// Get a connection that uses the native JDBC driver.  
  
Connection c = DriverManager.getConnection("jdbc:db2:*local");
```

The simplest form of JDBC URL is a list of three values that are separated by colons. The first value in the list represents the protocol which is always jdbc for JDBC URLs. The second value is the subprotocol and db2 is used to specify the native JDBC driver. The third value is the system name to establish the connection to a specific system. There are two special values that can be used to connect to the local database. They are *LOCAL and localhost (both are case insensitive). A specific system name can also be provided as follows:

```
Connection c =  
    DriverManager.getConnection("jdbc:db2:rchasmop");
```

This creates a connection to the rchasmop system. If the system to which you are trying to connect is a remote system (for example, through the Distributed Relational Database Architecture), the system name from the relational database directory must be used.

Note: When not specified, the user ID and password currently used to sign in are also used to establish the connection to the database.

Properties: The DriverManager.getConnection method takes a single string URL indicated previously and is only one of the methods on DriverManager to obtain a Connection object. There is also another version of the DriverManager.getConnection method that takes a user ID and password. The following is an example of this version:

Example: DriverManager.getConnection method taking a user ID and password

Note: Read the Code example disclaimer for important legal information.

```
// Get a connection that uses the native JDBC driver.  
  
Connection c = DriverManager.getConnection("jdbc:db2:*local", "cujo", "newtiger");
```

The line of code attempts to connect to the local database as user cujo with password newtiger no matter who is running the application. There is also a version of the DriverManager.getConnection method that takes a java.util.Properties object to allow further customization. The following is an example:

Example: DriverManager.getConnection method taking a java.util.Properties object

Note: Read the Code example disclaimer for important legal information.

```
// Get a connection that uses the native JDBC driver.  
  
Properties prop = new java.util.Properties();  
prop.put("user", "cujo");  
prop.put("password", "newtiger");  
Connection c = DriverManager.getConnection("jdbc:db2:*local", prop);
```

The code is functionally equivalent to the version previously mentioned that passed the user ID and password as parameters.

Refer to Connection properties for a complete list of connection properties for the native JDBC driver.

URL properties: Another way to specify properties is to place them in a list on the URL object itself. Each property in the list is separated by a semi-colon and the list must be of the form property name=property value. This is just a shortcut and does not significantly change the way processing is performed as the following example shows:

Example: Specify URL properties

Note: Read the Code example disclaimer for important legal information.

```
// Get a connection that uses the native JDBC driver.  
  
Connection c = DriverManager.getConnection("jdbc:db2:*local;user=cujo;password=newtiger");
```

The code is again functionally equivalent to the examples mentioned previously.

If a property value is specified in both a properties object and on the URL object, the URL version takes precedence over the version in the properties object. The following is an example:

Example: URL properties

Note: Read the Code example disclaimer for important legal information.

```
// Get a connection that uses the native JDBC driver.  
Properties prop = new java.util.Properties();  
prop.put("user", "someone");  
prop.put("password","something");  
Connection c = DriverManager.getConnection("jdbc:db2:*local;user=cujo;password=newtiger",  
prop);
```

The example uses the user ID and password from the URL string instead of the version in the Properties object. This ends up being functionally equivalent to the code previously mentioned.

See the following examples for more information:

- Use native JDBC and Toolbox JDBC concurrently
- Access property
- Invalid user ID and password



Example: Use native JDBC and Toolbox JDBC concurrently:  This is an example of how to use the native JDBC connection and the Toolbox JDBC connection in a program.

Example: Use native JDBC and Toolbox JDBC concurrently

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////  
//  
// GetConnections example.  
//  
// This program demonstrates being able to use both JDBC drivers at  
// once in a program. Two Connection objects are created in this  
// program. One is a native JDBC connection and one is a Toolbox  
// JDBC connection.  
//  
// This technique is convenient because it allows you to use different  
// JDBC drivers for different tasks concurrently. For example, the  
// Toolbox JDBC driver is ideal for connecting to remote iSeries servers  
// and the native JDBC driver is faster for local connections.
```

```

// You can use the strengths of each driver concurrently in your
// application by writing code similar to this example.
//
//
// This source is an example of the IBM Developer for Java JDBC driver.
// IBM grants you a nonexclusive license to use this as an example
// from which you can generate similar function tailored to
// your own specific needs.
//
// This sample code is provided by IBM for illustrative purposes
// only. These examples have not been thoroughly tested under all
// conditions. IBM, therefore, cannot guarantee or imply
// reliability, serviceability, or function of these programs.
//
// All programs contained herein are provided to you "AS IS"
// without any warranties of any kind. The implied warranties of
// merchantability and fitness for a particular purpose are
// expressly disclaimed.
//
// IBM Developer Kit for Java
// (C) Copyright IBM Corp. 2001
// All rights reserved.
// US Government Users Restricted Rights -
// Use, duplication, or disclosure restricted
// by GSA ADP Schedule Contract with IBM Corp.
//
import java.sql.*;
import java.util.*;

public class GetConnections {

    public static void main(java.lang.String[] args)
    {
        // Verify input.
        if (args.length != 2) {
            System.out.println("Usage (CL command line): java GetConnections PARM(<user> <password>");
            System.out.println(" where <user> is a valid iSeries user ID");
            System.out.println(" and <password> is the password for that user ID");
            System.exit(0);
        }

        // Register both drivers.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            Class.forName("com.ibm.as400.access.AS400JDBCdriver");
        } catch (ClassNotFoundException cnf) {
            System.out.println("ERROR: One of the JDBC drivers did not load.");
            System.exit(0);
        }

        try {
            // Obtain a connection with each driver.
            Connection conn1 = DriverManager.getConnection("jdbc:db2://localhost", args[0], args[1]);
            Connection conn2 = DriverManager.getConnection("jdbc:as400://localhost", args[0], args[1]);

            // Verify that they are different.
            if (conn1 instanceof com.ibm.db2.jdbc.app.DB2Connection)
                System.out.println("conn1 is running under the native JDBC driver.");
            else
                System.out.println("There is something wrong with conn1.");

            if (conn2 instanceof com.ibm.as400.access.AS400JDBCCConnection)
                System.out.println("conn2 is running under the Toolbox JDBC driver.");
            else
                System.out.println("There is something wrong with conn2.");
        }
    }
}

```

```

        conn1.close();
        conn2.close();
    } catch (SQLException e) {
        System.out.println("ERROR: " + e.getMessage());
    }
}
}

```



Example: Access property: ➤ This is an example of how to use the Access property.

Example: Access property

Note: Read the Code example disclaimer for important legal information.

```

// Note: This program assumes directory cujosql exists.
import java.sql.*;
import javax.sql.*;
import javax.naming.*;

public class AccessPropertyTest {
    public String url = "jdbc:db2:*local";
    public Connection connection = null;

    public static void main(java.lang.String[] args)
        throws Exception
    {
        AccessPropertyTest test = new AccessPropertyTest();

        test.setup();

        test.run();
        test.cleanup();
    }

    /**
    Set up the DataSource used in the testing.
    **/
    public void setup()
        throws Exception
    {
        Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

        connection = DriverManager.getConnection(url);
        Statement s = connection.createStatement();
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.TEMP");
        } catch (SQLException e) { // Ignore it - it doesn't exist
        }

        try {
            String sql = "CREATE PROCEDURE CUJOSQL.TEMP "
                + " LANGUAGE SQL SPECIFIC CUJOSQL.TEMP "
                + " MYPROC: BEGIN"
                + "     RETURN 11;"
                + " END MYPROC";
            s.executeUpdate(sql);
        } catch (SQLException e) {
            // Ignore it - it exists.
        }
        s.executeUpdate("create table cujosql.temp (col1 char(10))");
        s.executeUpdate("insert into cujosql.temp values ('compare')");
        s.close();
    }
}

```

```

}

public void resetConnection(String property)
throws SQLException
{
    if (connection != null)
        connection.close();

    connection = DriverManager.getConnection(url + ";access=" + property);
}

public boolean canQuery() {
    Statement s = null;
    try {
        s = connection.createStatement();
        ResultSet rs = s.executeQuery("SELECT * FROM cujosql.temp");
        if (rs == null)
            return false;

        rs.next();

        if (rs.getString(1).equals("compare "))
            return true;

        return false;
    } catch (SQLException e) {
        // System.out.println("Exception: SQLState(" + e.getSQLState() + ") " + e + " (" + e.getErrorCode() + ")");
        return false;
    } finally {
        if (s != null) {
            try {
                s.close();
            } catch (Exception e) {
                // Ignore it.
            }
        }
    }
}

public boolean canUpdate() {
    Statement s = null;
    try {
        s = connection.createStatement();
        int count = s.executeUpdate("INSERT INTO CUJOSQL.TEMP VALUES('x')");
        if (count != 1)
            return false;

        return true;
    } catch (SQLException e) {
        //System.out.println("Exception: SQLState(" + e.getSQLState() + ") " + e + " (" + e.getErrorCode() + ")");
        return false;
    } finally {
        if (s != null) {
            try {
                s.close();
            } catch (Exception e) {
                // Ignore it.
            }
        }
    }
}
}

```

```

public boolean canCall() {
    CallableStatement s = null;
    try {
        s = connection.prepareCall("? = CALL CUJOSQL.TEMP()");
        s.registerOutParameter(1, Types.INTEGER);
        s.execute();
        if (s.getInt(1) != 11)
            return false;

        return true;
    } catch (SQLException e) {
        //System.out.println("Exception: SQLState(" + e.getSQLState() + ") " + e + " (" + e.getErrorCode() + ")");
        return false;
    } finally {
        if (s != null) {
            try {
                s.close();
            } catch (Exception e) {
                // Ignore it.
            }
        }
    }
}

```

```

public void run()
throws SQLException
{
    System.out.println("Set the connection access property to read only");
    resetConnection("read only");

    System.out.println("Can run queries -->" + canQuery());
    System.out.println("Can run updates -->" + canUpdate());
    System.out.println("Can run sp calls -->" + canCall());

    System.out.println("Set the connection access property to read call");
    resetConnection("read call");

    System.out.println("Can run queries -->" + canQuery());
    System.out.println("Can run updates -->" + canUpdate());
    System.out.println("Can run sp calls -->" + canCall());

    System.out.println("Set the connection access property to all");
    resetConnection("all");

    System.out.println("Can run queries -->" + canQuery());
    System.out.println("Can run updates -->" + canUpdate());
    System.out.println("Can run sp calls -->" + canCall());
}

```

```

public void cleanup() {
    try {
        connection.close();
    } catch (Exception e) {
        // Ignore it.
    }
}
}

```



Example: Invalid user ID and password: ➤ This is an example of how to use the Connection property in SQL naming mode.

Example: Invalid user ID and password

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
//
// InvalidConnect example.
//
// This program uses the Connection property in SQL naming mode.
//
////////////////////////////////////
//
// This source is an example of the IBM Developer for Java JDBC driver.
// IBM grants you a nonexclusive license to use this as an example
// from which you can generate similar function tailored to
// your own specific needs.
//
// This sample code is provided by IBM for illustrative purposes
// only. These examples have not been thoroughly tested under all
// conditions. IBM, therefore, cannot guarantee or imply
// reliability, serviceability, or function of these programs.
//
// All programs contained herein are provided to you "AS IS"
// without any warranties of any kind. The implied warranties of
// merchantability and fitness for a particular purpose are
// expressly disclaimed.
//
// IBM Developer Kit for Java
// (C) Copyright IBM Corp. 2001
// All rights reserved.
// US Government Users Restricted Rights -
// Use, duplication, or disclosure restricted
// by GSA ADP Schedule Contract with IBM Corp.
//
////////////////////////////////////
import java.sql.*;
import java.util.*;

public class InvalidConnect {

    public static void main(java.lang.String[] args)
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (ClassNotFoundException cnf) {
            System.out.println("ERROR: JDBC driver did not load.");
            System.exit(0);
        }

        // Attempt to obtain a connection without specifying any user or
        // password. The attempt works and the connection uses the
        // same user profile under which the job is running.
        try {
            Connection c1 = DriverManager.getConnection("jdbc:db2:*local");
            c1.close();
        } catch (SQLException e) {
            System.out.println("This test should not get into this exception path.");
            e.printStackTrace();
            System.exit(1);
        }

        try {
```

```



        Connection c2 = DriverManager.getConnection("jdbc:db2:*local", "notvalid", "notvalid");
    } catch (SQLException e) {
        System.out.println("This is an expected error.");
        System.out.println("Message is " + e.getMessage());
        System.out.println("SQLSTATE is " + e.getSQLState());
    }
}
}

```



Connection properties: This table contains valid JDBC driver connection properties, their values, and their descriptions:

Property	Values	Meaning
"access"	"all", "read call", "read only"	This value can be used to restrict the type of operations that can be done with a specific connection. The default value is "all" and basically means that the connection has full access to the JDBC API. The "read call" value allows the connection to do only queries and call stored procedures. An attempt to update the database through an SQL statement is stopped. The "read only" value can be used to restrict a connection to only queries. Stored procedure calls and update statements are stopped.
➤ "batch style"	"2.0", "2.1"	The JDBC 2.1 specification defines a second method for how exceptions in a batch update can be handled. The driver can comply with either of these. The default is to work as defined in the JDBC 2.0 specification. ⬅

Property	Values	Meaning
"block size"	"0", "8", "16", "32", "64", "128", "256", "512"	<p>This is the number of rows that are fetched at a time for a result set. For typical forward only processing of a result set, a block of this size are obtained. Then the database is not accessed as each row is processed by your application. Only when the end of the block is reached that the database requests another block of data.</p> <p>This value is only used if the "blocking enabled" property is set to true.</p> <p>Setting the block size property to "0" has the same effect as setting the "blocking enabled" property to "false".</p> <p>The default is to use blocking with a block size of "32". Currently, this is a fairly arbitrary decision and the default could change in the future.</p> <p>Currently, blocking is not used on scrollable result sets.</p>
"blocking enabled"	"true", "false"	<p>This value is used to determine whether or not the connection should use blocking on result set row retrieval. Blocking can significantly improve the performance of processing result sets.</p> <p>By default, this property is set to true.</p>
 "cursor hold"	"true", "false"	<p>This value specifies whether or not result sets should remain open when a transaction is committed. A value of true means that an application is able to access its open result sets after commit is called. A value of false means that commit closes any open cursors under the connection.</p> <p>By default, this property is set to true.</p> <p>This value property serves as a default value for all result sets made for the connection. With cursor holdability support added in JDBC 3.0, this default is simply replaced if an application specifies a different holdability later. </p>

Property	Values	Meaning
» "data truncation"	"true", "false"	This value specifies whether truncation of character data should cause warnings and exceptions to be generated (true) or if the data should just be silently truncated (false). If the default is true, data truncation of character fields should be honored. <<
» "date format"	"julian", "mdy", "dmy", "ymd", "usa", "iso", "eur", "jis"	This property allows you to change how dates are formatted. <<
» "date separator"	"/", "-", ".", " ", "b"	This property allows you to change what the date separator is. This is only valid in combination with some of the dateFormat values (according to system rules). <<
» "decimal separator"	(".", " ")	This property allows you to change what the decimal separator is. <<

Property	Values	Meaning
"do escape processing"	"true", "false"	<p>This property sets a flag for whether or not statements under the connection must do escape processing. Using escape processing is a way to code your SQL statements so that they are generic and similar for all platforms, but then the database reads the escape clauses and substitutes the proper system specific version for the user.</p> <p>This is good, except that it forces extra work on the system. In the case where you know you are only using SQL statements that already contain valid iSeries SQL syntax, it is recommended that this value be set to "false" to increase performance.</p> <p>The default value for this property is true, as it must be compliant with the JDBC specification (that is, escape processing is active by default).</p> <p>This value is added due to a shortcoming of the JDBC specification. You can only set escape processing to off in the Statement class. That works well if you are dealing with simple statements. You create your statement, turn off escape processing, and start processing statements. However, in the case of prepared statements and callable statements, this scheme does not work. You supply the SQL statement at the time that the prepared or callable statement is constructed and it does not change after that. So the statement is prepared up front and changing the escape processing after that is meaningless. Having this connection property allows a way to get around the extra overhead.</p>
 "errors"	"basic", "full"	<p>This property allows the full system second-level error text to be returned in SQLException object messages. The default is basic which returns only the standard message text. </p>

Property	Values	Meaning
<p>» "libraries"</p>	<p>A space-separated list of libraries. (A list of libraries can also be separated by colons or commas.)</p>	<p>This property allows a list of libraries to be placed into the server job's library list or a specific default library to be set.</p> <p>The "naming" property affects how this property works. In the default case, where "naming" is set to "sql", JDBC works like ODBC. The library list has no effect on how the connection processes. There is a default library for all unqualified tables. By default, that library has the same name as the user profile that is connected. If the libraries property is specified, the first library in the list becomes the default library. If a default library is specified on the connection URL (as in "jdbc:db2:*local/mylibrary"), that overrides any value in this property.</p> <p>In the case where "naming" is set "system", each of the libraries specified for this property is added to the user's library list and the library list is searched to resolve unqualified table references. <<</p>
<p>» "lob threshold"</p>	<p>Any value under 500000</p>	<p>This property tells the driver to place the actual values into the result set storage instead of locators for lob columns if the lob column is smaller than the threshold size. This property acts against the column size, not the lob data size itself. For example, if the lob column is defined to hold up to 1 MB for each lob, but all the column values are under 500 KB, locators are still used.</p> <p>Note that the size limit is set as it is to allow blocks of data to be fetched without danger of not always growing data blocks larger than the 16 MB maximum allocation size. With large result sets, it is still easy to exceed this limit, which causes fetches to fail. Care must be taken in how the block size property and this property interact with the size of a data block.</p> <p>The default is 0. Locators are always used for lob data. <<</p>

Property	Values	Meaning
"naming"	"sql", "system"	<p>This property allows you to use either the traditional iSeries naming syntax or the standard sql naming syntax. "system" naming means that you should use a "/" character to separate collection and table values, and "sql" naming means that you should use a "." character to separate the values.</p> <p>The setting of this value has ramifications for what the default library is also. See the libraries property for further information on this.</p> <p>The default is to use "sql" naming.</p>
"password"	anything	<p>This property allows for a password to be specified for the connection. This property does not work correctly without also specifying the "user" property. These properties allow for connections to be made to the database as a user other than the one that is running the iSeries job.</p> <p>Specifying the "user" and "password" properties have the same effect as using the connection method with the signature getConnection(String url, String userId, String password).</p>
» "prefetch"	"true", "false"	<p>This property specifies whether or not the driver should fetch the first data for a result set immediately after processing or wait until the data is requested. If the default is true, data should be prefetched.</p> <p>For applications using the Native JDBC driver, this is rarely an issue. The property exists primarily for internal use with Java stored procedures and user-defined functions where it is important that the database engine does not fetch any data from result sets on your behalf before you request it. «</p>
» "reuse objects"	"true", "false"	<p>This property specifies whether or not the driver should attempt to reuse some types of objects after you close them. This is a performance enhancement. The default is true. «</p>

Property	Values	Meaning
» "server trace"	A string representation of an integer	<p>This property enables tracing of the JDBC server job. If server tracing is enabled, tracing starts when the client connects to the server, and ends when the connection is disconnected.</p> <p>Trace data is collected in spooled files on the server. Multiple levels of server tracing can be turned on in combination by adding the constants and passing that sum on the set method.</p> <p>Note: This property is typically used by support personnel and its values are not discussed further. <<</p>
» "time format"	"hms", "usa", "iso", "eur", "jis"	This property allows you to change how time values are formatted. <<
» "time separator"	":", ".", ",", "b"	This property allows you to change what the time separator is. This is only valid in combination with some of the timeFormat values (according to system rules). <<
"trace"	"true", "false"	<p>This property allows for turning on tracing of the connection. It can be used as a simple debugging aide. There are thoughts of enhancing this feature in the future. See D2. The JDBC driver threw an exception. What do I do? for more information on debugging.</p> <p>The default value is "false", which does not use tracing.</p>
"transaction isolation"	"none", "read committed", "read uncommitted", "repeatable read", "serializable"	<p>This property allows you to set the transaction isolation level for the connection. There is no difference between setting this property to a specific level and specifying a level on the setTransactionIsolation method in the Connection interface.</p> <p>The default value for this property is "none", as JDBC defaults to auto-commit mode.</p>
"translate binary"	"true", "false"	<p>This property can be used to force the JDBC driver to treat binary and varbinary data values as if they were char and varchar data values.</p> <p>The default for this property is "false", where binary data is not treated the same as character data.</p>

Property	Values	Meaning
<p>» "use block insert"</p>	<p>"true", "false"</p>	<p>This property allows the native JDBC driver to go into a block insert mode for inserting blocks of data into the database. This is an optimized version of the batch update. This optimized mode can only be used in applications that ensure that they do not break certain system constraints or data insert failures and potentially corrupt data.</p> <p>Applications that turn on this property should only connect to the local system when attempting to perform batched updates. They should not use DRDA to establish remote connections because blocked insert cannot be managed over DRDA.</p> <p>Applications must also ensure that PreparedStatements with an SQL insert statement and a values clause make all the insert values parameters. No constants are permitted in the values list. This is a requirement of the blocked insert engine of the system.</p> <p>The default is false. «</p>
<p>"user"</p>	<p>anything</p>	<p>This property allows for a user ID to be specified for the connection. This property does not work correctly without also specifying the "password" property. These properties allow for connections to be made to the database as a user other than the one that is running the iSeries job.</p> <p>Specifying the "user" and "password" properties has the same effect as using the connection method with the signature getConnection(String url, String userId, String password).</p>

Example: Create a UDBDataSource and bind it with JNDI: » This is an example of how to create a UDBDataSource and get it bound with JNDI.

Example: Create a UDBDataSource and bind it with JNDI

Note: Read the Code example disclaimer for important legal information.

```
// Import the required packages. At deployment time,
// the JDBC driver-specific class that implements
// DataSource must be imported.
import java.sql.*;
import javax.naming.*;
import com.ibm.db2.jdbc.app.UDBDataSource;
```

```

public class UDBDataSourceBind
{
    public static void main(java.lang.String[] args)
    throws Exception
    {
        // Create a new UDBDataSource object and give it
        // a description.
        UDBDataSource ds = new UDBDataSource();
        ds.setDescription("A simple UDBDataSource");

        // Retrieve a JNDI context. The context serves
        // as the root for where objects are bound or
        // found in JNDI.
        Context ctx = new InitialContext();

        // Bind the newly created UDBDataSource object
        // to the JNDI directory service, giving it a name
        // that can be used to look up this object again
        // at a later time.
        ctx.rebind("SimpleDS", ds);
    }
}

```

Example: Create a UDBDataSourceBind and set DataSource properties: ➤ This is an example of how to create a UDBDataSource, and set the user ID and password as DataSource properties.

Example: Create a UDBDataSourceBind and set DataSource properties

Note: Read the Code example disclaimer for important legal information.

```

// Import the required packages. At deployment time,
// the JDBC driver-specific class that implements
// DataSource must be imported.
import java.sql.*;
import javax.naming.*;
import com.ibm.db2.jdbc.app.UDBDataSource;

public class UDBDataSourceBind2
{
    public static void main(java.lang.String[] args)
    throws Exception
    {
        // Create a new UDBDataSource object and give it
        // a description.
        UDBDataSource ds = new UDBDataSource();
        ds.setDescription("A simple UDBDataSource " +
            "with cujo as the default " +
            "profile to connect with.");

        // Provide a user ID and password to be used for
        // connection requests.
        ds.setUser("cujo");
        ds.setPassword("newtiger");

        // Retrieve a JNDI context. The context serves
        // as the root for where objects are bound or
        // found in JNDI.
        Context ctx = new InitialContext();

        // Bind the newly created UDBDataSource object
        // to the JNDI directory service, giving it a name
        // that can be used to look up this object again
        // at a later time.
        ctx.rebind("SimpleDS2", ds);
    }
}

```

Example: Obtain an initial context before binding UDBDataSource: ➤ The following example obtains an initial context before binding the UDBDataSource. The lookup method is then used on that context to return an object of type DataSource for the application to use.

Example: Obtain an initial context before binding UDBDataSource

Note: Read the Code example disclaimer for important legal information.

```
// Import the required packages. There is no
// driver-specific code needed in runtime
// applications.
import java.sql.*;
import javax.sql.*;
import javax.naming.*;

public class UDBDataSourceUse
{
    public static void main(java.lang.String[] args)
        throws Exception
    {
        // Retrieve a JNDI context. The context serves
        // as the root for where objects are bound or
        // found in JNDI.
        Context ctx = new InitialContext();

        // Retrieve the bound UDBDataSource object using the
        // name with which it was previously bound. At runtime,
        // only the DataSource interface is used, so there
        // is no need to convert the object to the UDBDataSource
        // implementation class. (There is no need to know what
        // the implementation class is. The logical JNDI name is
        // only required).
        DataSource ds = (DataSource) ctx.lookup("SimpleDS");

        // Once the DataSource is obtained, it can be used to establish
        // a connection. This Connection object is the same type
        // of object that is returned if the DriverManager approach
        // to establishing connection is used. Thus, so everything from
        // this point forward is exactly like any other JDBC
        // application.
        Connection connection = ds.getConnection();

        // The connection can be used to create Statement objects and
        // update the database or process queries as follows.
        Statement statement = connection.createStatement();
        ResultSet rs = statement.executeQuery("select * from qsys2.sysprocs");
        while (rs.next()) {
            System.out.println(rs.getString(1) + "." + rs.getString(2));
        }

        // The connection is closed before the application ends.
        connection.close();
    }
}
```

Example: Create a UDBDataSource, and obtain a user ID and password: ➤ This is an example of how to create a UDBDataSource, and use the getConnection method to obtain a user ID and password at runtime.

Example: Create a UDBDataSource, and obtain a user ID and password

Note: Read the Code example disclaimer for important legal information.

```

/// Import the required packages. There is
// no driver-specific code needed in runtime
// applications.
import java.sql.*;
import javax.sql.*;
import javax.naming.*;

public class UDBDataSourceUse2
{
    public static void main(java.lang.String[] args)
    throws Exception
    {
        // Retrieve a JNDI context. The context serves
        // as the root for where objects are bound or
        // found in JNDI.
        Context ctx = new InitialContext();

        // Retrieve the bound UDBDataSource object using the
        // name with which it was previously bound. At runtime,
        // only the DataSource interface is used, so there
        // is no need to convert the object to the UDBDataSource
        // implementation class. (There is no need to know
        // what the implementation class is. The logical JNDI name
        // is only required).
        DataSource ds = (DataSource) ctx.lookup("SimpleDS");

        // Once the DataSource is obtained, it can be used to establish
        // a connection. The user profile cujo and password newtiger
        // used to create the connection instead of any default user
        // ID and password for the DataSource.
        Connection connection = ds.getConnection("cujo", "newtiger");

        // The connection can be used to create Statement objects and
        // update the database or process queries as follows.
        Statement statement = connection.createStatement();
        ResultSet rs = statement.executeQuery("select * from qsys2.sysprocs");
        while (rs.next()) {
            System.out.println(rs.getString(1) + "." + rs.getString(2));
        }

        // The connection is closed before the application ends.
        connection.close();
    }
}

```

Use DataSources with UDBDataSource: ➤ DataSource interfaces were designed to allow additional flexibility in using Java^(TM) Database Connectivity (JDBC) drivers. The use of DataSources can be split into two phases:

- **Deployment**


Deployment is a setup phase that occurs before a JDBC application actually runs. Deployment usually involves setting up a DataSource to have specific properties and then binding it into a directory service through the use of the Java Naming and Directory Interface (JNDI). The directory service is most commonly the Lightweight Directory Access Protocol (LDAP), but could be a number of others such as Common Object Request Broker Architecture (CORBA) Object Services, Java Remote Method Invocation (RMI), or the underlying file system.

- **Use**

By decoupling the deployment from the runtime use of the DataSource, the DataSource setup can be reused by many applications. By changing some aspect of the deployment, all the applications that use that DataSource automatically pick up the changes.

An advantage of DataSources is that they allow JDBC drivers to do work on behalf of the application without having an impact on the application development process directly. For more information, see Connection pooling, Statement pooling, and Distributed transactions.

UDBDataSourceBind: The `UDBDataSourceBind` program is an example of creating a `UDBDataSource` and getting it bound with JNDI. This program accomplishes all the basic tasks requested. Namely, it instantiates a `UDBDataSource` object, sets properties on this object, retrieves a JNDI context, and binds the object to a name within the JNDI context.

The deployment time code is vendor-specific. The application must import the specific `DataSource` implementation that it wants to work with. In the import list, the package-qualified `UDBDataSource` class is imported. The most unfamiliar part of this application is the work done with JNDI (for example, the retrieval of the Context object and the call to bind). For additional information, see [JNDI](#)  by Sun Microsystems, Inc.

Once this program has been run and has successfully completed, there is a new entry in a JNDI directory service called SimpleDS. This entry is at the location specified by the JNDI context. The `DataSource` implementation is now deployed. An application program can make use of this `DataSource` to retrieve database connections and JDBC-related work.

UDBDataSourceUse: The `UDBDataSourceUse` program is an example of a JDBC application that uses the previously deployed application.

The JDBC application obtains an initial context as it did before binding the `UDBDataSource` in the previous example. The lookup method is then used on that context to return an object of type `DataSource` for the application to use.

Note: The runtime application is only interested in the methods of the `DataSource` interface, so there is no need for it to be aware of the implementation class. This makes applications portable.

Suppose that `UDBDataSourceUse` is a complex application that runs a large operation within your organization. You have a dozen or more similar large applications within your organization. You have to change the name of one of the systems in your network. By running a deployment tool and changing a single `UDBDataSource` property, you would be able to get this new behavior in all your applications without changing the code for them. One of the benefits of `DataSources` is that they allow you to consolidate system setup information. Another major benefit is that they allow drivers to implement functionality invisible to the application such as connection pooling, statement pooling and support for distributed transactions.

After analyzing `UDBDataSourceBind` and `UDBDataSourceUse` closely, you may have wondered how the `DataSource` object knew what to do. There is no code to specify a system, a user ID, or a password in either of these programs. The `UDBDataSource` class has default values for all properties; by default, it connects to the local iSeries server with the user profile and password of the running application. If you wanted to ensure that the connection was made with the user profile `cujo` instead, you could have accomplished this in two ways:

- Set the user ID and password as `DataSource` properties. See [Example: Create a UDBDataSourceBind and set DataSource properties](#) on how to use this technique.
- Use the `DataSource` `getConnection` method that takes a user ID and password at runtime. See [Example: Create a UDBDataSource, and obtain a user ID and password](#) on how to use this technique.


There are a number of properties that can be specified for the `UDBDataSource` as there are properties that can be specified for connections created with the `DriverManager`. Refer to `DataSource` properties for a list of supported properties for the native JDBC driver.

While these lists are similar, it is not certain to be similar in future releases. You are encouraged to start coding to the DataSource interface.

Note: The native JDBC driver also has two other DataSource implementations, but direct use of them is not recommended.

- DB2DataSource
- DB2StdDataSource



DataSource properties:  This table contains valid data source properties, their values, and their descriptions:

Set method (data type)	Values	Description
setAccess (String)	"all", "read call", "read only"	<p>This property can be used to restrict the type of operations that can be done with a specific connection. The default value is "all" and basically means that the connection has full access to the Java^(TM) Database Connectivity (JDBC) API.</p> <p>The "read call" value allows the connection to only perform queries and call stored procedures. An attempt to update the database through an SQL statement causes an SQLException.</p> <p>The "read only" value restricts the connection to only queries. An attempt to process a stored procedure call or update statements causes an SQLException.</p>
setBatchStyle (String)	"2.0", "2.1"	<p>The JDBC 2.1 specification defines a second method for how exceptions in a batch update can be handled. The driver can comply with either of these. The default is to work as defined in the JDBC 2.0 specification.</p>
setUseBlocking (boolean)	"true", "false"	<p>This property is used to determine whether or not the connection should use blocking on result set row retrieval. Blocking can significantly improve the performance of processing result sets.</p> <p>By default, this property is set to true.</p>

Set method (data type)	Values	Description
<p>setBlockSize (int)</p>	<p>"0", "8", "16", "32", "64", "128", "256", "512"</p>	<p>This property indicates the number of rows that are fetched at a time for a result set. For typical forward only processing of a result set, a block of this size is obtained if the database has that many rows that satisfy the query. Only when the end of the block is reached in the JDBC driver's internal storage, another request for a block of data is sent to the database.</p> <p>This value is only used if the useBlocking property is set to true. Refer to setUseBlocking for more information.</p> <p>Setting the block size property to "0" has the same effect as calling setUseBlocking(false).</p> <p>The default is to use blocking with a block size of "32". This is a fairly arbitrary decision and the default could change in future releases.</p> <p>Currently, blocking is not used on scrollable result sets.</p> <p>Using blocking affects the degree of cursor sensitivity the user application has. A sensitive cursor sees changes made by other SQL statements. However, because of data caching, changes are only detected when data needs to be fetched from the database.</p>
<p>setCursorHold (boolean)</p>	<p>"true", "false"</p>	<p>This property specifies whether or not result sets should remain open when a transaction is committed. A value of true means that an application is able to access its open result sets after commit is called. A value of false means that commit closes any open cursors under the connection.</p> <p>By default, this property is set to true.</p> <p>This property serves as a default value for all result sets made for the connection. With cursor support added in JDBC 3.0 (see the ResultSet characteristics section for details), this default is simply replaced if an application specifies different cursor support later.</p>

Set method (data type)	Values	Description
setDataTruncation (boolean)	"true", "false"	This property specifies the following: <ul style="list-style-type: none"> Whether truncation of character data should cause warnings and exceptions to be generated (true) If the data should just be silently truncated (false). See DataTruncation for additional details.
setDatabaseName (String)	Any name	This property specifies the database to which the DataSource attempts to connect. The default is *LOCAL. The database name must either exist in the relational database directory on the system that runs the application or be the special value *LOCAL or localhost to specify the local system.
setDataSourceName (String)	Any name	This property allows the passing of a ConnectionPoolDataSource Java Naming and Directory Interface (JNDI) name to support connection pooling.
setDateFormat (String)	"julian", "mdy", "dmy", "ymd", "usa", "iso", "eur", "jis"	This property allows you to change how dates are formatted.
setDateSeparator (String)	"/", "-", ".", ",", "b"	This property allows you to change what the date separator is. This is only valid in combination with some of the dateFormat values (according to system rules).
setDecimalSeparator (String)	".", ",", ""	This property allows you to change what the decimal separator should be.
setDescription (String)	Any name	This property allows the setting of this DataSource object's text description.
setDoEscapeProcessing (boolean)	"true", "false"	This property specifies whether SQL statements have escaped processing done on them. The default value for this property is true.
setFullErrors (boolean)	"true", "false"	This property allows second-level error text of the full system to be returned in SQLException object messages. The default is false.
setLibraries (String)	A space-separated list of libraries	This property allows a list of libraries to be placed into the server job's library list. This property is only used when setSystemNaming(true) is used.
setLobThreshold (int)	Any value under 500000	This property tells the driver to place the actual values instead of Locator Object (LOB) locators if the LOB column is smaller than the threshold size.

Set method (data type)	Values	Description
setLoginTimeout (int)	Any value	This property is currently ignored and is planned for future use.
setNetworkProtocol (int)	Any value	This property is currently ignored and is planned for future use.
setPassword (String)	Any string	This property allows for a password to be specified for the connection. This property is ignored if a user ID is not set.
setPortNumber (int)	Any value	This property is currently ignored and is planned for future use.
setPrefetch (boolean)	"true", "false"	This property specifies whether the driver should fetch the first data for a result set immediately after processing or wait until the data is requested. The default is true.
setReuseObjects (boolean)	"true", "false"	This property specifies whether the driver should attempt to reuse some types of objects after you close them. This is a performance enhancement. The default is true.
setServerName (String)	Any name	This property is currently ignored and is planned for future use.
setServerTraceCategories (int)	A string representation of an integer	<p>This property enables tracing of the JDBC server job. If server tracing is enabled, tracing starts when the client connects to the server, and ends when the connection is disconnected.</p> <p>Trace data is collected in spooled files on the server. Multiple levels of server tracing can be turned on in combination by adding the constants and passing that sum on the set method.</p> <p>Note: This property is typically used by support personnel and its values are not discussed further.</p>
setSystemNaming (boolean)	"true", "false"	This property allows specifying whether collections and tables should be separated by a period (SQL naming) or a slash (system naming). This property also determines whether a default library is used (SQL naming) or the library list is used (system naming). The default is SQL naming.
setTimeFormat (String)	"hms", "usa", "iso", "eur", "jis"	This property allows you to change how time values are formatted.
setTimeSeparator (String)	".", ".", ",", "b"	This property allows you to change what the time separator is. This is only valid in combination with some of the timeFormat values (according to system rules).

Set method (data type)	Values	Description
setTrace (boolean)	"true", "false"	This property can enable a simple trace. The default value is false.
setTransactionIsolationLevel (String)	"none", "read committed", "read uncommitted", "repeatable read", "serializable"	This property allows the specification of the transaction isolation level. The default value for this property is "none", as JDBC defaults to auto-commit mode.
setTranslateBinary (Boolean)	"true", "false"	This property can be used to force the JDBC driver to treat binary and varbinary data values as if they were char and varchar data values. The default for this property is false.
setUseBlockInsert (boolean)	"true", "false"	This property allows the native JDBC driver to go into a block insert mode for inserting blocks of data into the database. This is an optimized version of the batch update. This optimized mode can only be used in applications that ensure that they do not break certain system constraints or data insert failures and potentially corrupt data. Applications that turn on this property should only connect to the local system when attempting to perform batched updates. They should not use DRDA to establish remote connections because a blocked insert cannot be managed over DRDA. Applications must also ensure that PreparedStatements with an SQL insert statement and a values clause make all the insert values parameters. No constants are permitted in the values list. This is a requirement of the blocked insert engine of the system. The default is false.
setUser (String)	anything	This property allows the setting of a user ID for obtaining connections. This property requires that you also set the password property.



Other DataSource implementations: ➤ There are two implementations of the DataSource interface that are included with the native JDBC driver. These DataSource implementations should be considered deprecated. While you can still use them, they are not enhanced with future improvements; for example, robust connection and statement pooling are not added to these implementations. These implementations exist until you adopt the UDBDataSource interface and its related functions.

DB2DataSource: The *DB2DataSource* was an early implementation of the *DataSource* interface and does not comply with the complete specification (that is, it predates the specification). *DB2DataSource* exists today only to allow WebSphere^(R) users to migrate to current releases and should not be used otherwise.

DB2StdDataSource: The *DB2StdDataSource* is the revised version of the *DB2DataSource* implementation that became specification-compliant once the JDBC optional package specification became final. The new version was provided to not break code already written on the *DB2DataSource* version.

If you have written applications that make use of these *DataSource* implementations, migrating to the *UDBDataSource* is a trivial task as all the old properties are supported. It is recommended that you migrate to *UDBDataSource* to gain the functionality of the new *UDBDataSource* classes. <<

DatabaseMetaData interface for IBM Developer Kit for Java

The *DatabaseMetaData* interface is implemented by the IBM Developer Kit for Java^(TM) JDBC driver to provide information about its underlying data sources. It is used primarily by application servers and tools to determine how to interact with a given data source. Applications may also use *DatabaseMetaData* methods to obtain information about a data source, but this is less typical.

The *DatabaseMetaData* interface includes over 150 methods that can be categorized according to the following types of information they provide:

- General information about the data source
- Data source support for a given feature
- Data source limits
- SQL objects and their attributes
- Transaction support offered by the data source

The *DatabaseMetaData* interface also contains over 40 fields that are constants used as return values for various *DatabaseMetaData* methods.

See Changes in JDBC 3.0 for information about changes made to methods in the *DatabaseMetaData* interface.

Create a DatabaseMetaData object: A *DatabaseMetaData* object is created with the *Connection* method *getMetaData*. Once the object is created, it can be used to dynamically find information about the underlying data source. The following example creates a *DatabaseMetaData* object and uses it to determine the maximum number of characters allowed for a table name:

Example: Create a *DatabaseMetaData* object

Note: Read the Code example disclaimer for important legal information.

```
// con is a Connection object
DatabaseMetaData dbmd = con.getMetadadata();
int maxLen = dbmd.getMaxTableNameLength();
```

Retrieve general information: Some *DatabaseMetaData* methods are used to dynamically find general information about a data source as well as to obtain details about its implementation. Some of these methods include the following:

- *getURL*
- *getUserName*
- *getDatabaseProductVersion*, *getDriverMajorVersion*, and *getDriverMinorVersion*
- *getSchemaTerm*, *getCatalogTerm*, and *getProcedureTerm*
- *nullsAreSortedHigh*, and *nullsAreSortedLow*
- *usesLocalFiles*, and *usesLocalFilePerTable*

- getSQLKeywords

Determine feature support: A large group of DatabaseMetaData methods can be used to determine whether a given feature or set of features is supported by the driver or underlying data source. Beyond this, there are methods that describe what level of support is provided. Some of the methods that describe support for individual features include the following:

- supportsAlterTableWithDropColumn
- supportsBatchUpdates
- supportsTableCorrelationNames
- supportsPositionedDelete
- supportsFullOuterJoins
- supportsStoredProcedures
- supportsMixedCaseQuotedIdentifiers

Methods to describe a level of feature support include the following:

- supportsANSI92EntryLevelSQL
- supportsCoreSQLGrammar

Data source limits: Another group of methods provide the limits imposed by a given data source. Some of the methods in this category include the following:

- getMaxRowSize
- getMaxStatementLength
- getMaxTablesInSelect
- getMaxConnections
- getMaxCharLiteralLength
- getMaxColumnsInTable

Methods in this group return the limit value as an integer. A return value of zero means that there is either no limit or the limit is unknown.


SQL objects and their attributes: A number of DatabaseMetaData methods provide information about the SQL objects that populate a given data source. These methods can determine the attributes of SQL objects. These methods also return ResultSet objects in which each row describes a particular object. For example, the getUDTs method returns a ResultSet object in which there is a row for each user-defined table (UDT) that has been defined in the data source. Examples of this category include the following:

- getSchemas and getCatalogs
- getTables
- getPrimaryKeys
- getProcedures and getProcedureColumns
- getUDTs


Transaction support: A small group of methods provide information about the transaction semantics supported by the data source. Examples of this category include the following:

- supportsMultipleTransactions
- getDefaultTransactionIsolation

See Example: DatabaseMetaData interface for IBM Developer Kit for Java for an example of how to use the DatabaseMetaData interface.

Changes in JDBC 3.0:  There are changes to the return values for some of the methods in JDBC 3.0. The following methods have been updated in JDBC 3.0 to add fields to the ResultSets that they return.

- getTables
- getColumns
- getUDTs
- getSchemas

Note: If an application is being developed using Java Development Kit (JDK) 1.4, you may recognize that there are a certain number of columns being returned when testing. You write your application and expect to access all of these columns. However, if the application is being designed to also run on previous releases of the JDK, the application receives an SQLException when it tries to access these fields that do not exist in earlier JDK releases. SafeGetUDTs is an example of how an application can be written to work with JDK 1.4, JDK 1.3, and previous JDK releases. 

Example: DatabaseMetaData interface for IBM Developer Kit for Java: This example shows how to return a list of tables.

Example 1: Return a list of tables

Note: Read the Code example disclaimer for important legal information.

```
// Connect to iSeries server.
Connection c = DriverManager.getConnection("jdbc:db2:mySystem");


// Get the database meta data from the connection.
DatabaseMetaData dbMeta = c.getMetaData();

// Get a list of tables matching this criteria.
String catalog = "myCatalog";
String schema = "mySchema";
String table = "myTable%"; // % indicates search pattern
String types[] = {"TABLE", "VIEW", "SYSTEM TABLE"};
ResultSet rs = dbMeta.getTables(catalog, schema, table, types);

// ... iterate through the ResultSet to get the values.

// Close the connection.
c.close();
```

For more information, see DatabaseMetaData interface for IBM Developer Kit for Java^(TM).

Example: Use metadata ResultSets that have more than one column:  This is an example of how to use metadata ResultSets that have more than one column.

Example: Use metadata ResultSets that have more than one column

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
//
// SafeGetUDTs example. This program demonstrates one way to deal with
// metadata ResultSets that have more columns in JDK 1.4 than they
// had in previous releases.
//
// Command syntax:
//   java SafeGetUDTs
//
////////////////////////////////////
//
// This source is an example of the IBM Developer for Java JDBC driver.
// IBM grants you a nonexclusive license to use this as an example
// from which you can generate similar function tailored to
// your own specific needs.
//
```

```

// This sample code is provided by IBM for illustrative purposes
// only. These examples have not been thoroughly tested under all
// conditions. IBM, therefore, cannot guarantee or imply
// reliability, serviceability, or function of these programs.
//
// All programs contained herein are provided to you "AS IS"
// without any warranties of any kind. The implied warranties of
// merchantability and fitness for a particular purpose are
// expressly disclaimed.
//
// IBM Developer Kit for Java
// (C) Copyright IBM Corp. 2001
// All rights reserved.
// US Government Users Restricted Rights -
// Use, duplication, or disclosure restricted
// by GSA ADP Schedule Contract with IBM Corp.
//
////////////////////////////////////

import java.sql.*;

public class SafeGetUDTs {

    public static int jdbcLevel;

    // Note: Static block runs before main begins.
    // Therefore, there is access to jdbcLevel in
    // main.
    {
        try {
            Class.forName("java.sql.Blob");

            try {
                Class.forName("java.sql.ParameterMetaData");
                // Found a JDBC 3.0 interface. Must support JDBC 3.0.
                jdbcLevel = 3;
            } catch (ClassNotFoundException ez) {
                // Could not find the JDBC 3.0 ParameterMetaData class.
                // Must be running under a JVM with only JDBC 2.0
                // support.
                jdbcLevel = 2;
            }

        } catch (ClassNotFoundException ex) {
            // Could not find the JDBC 2.0 Blob class. Must be
            // running under a JVM with only JDBC 1.0 support.
            jdbcLevel = 1;
        }
    }

    // Program entry point.
    public static void main(java.lang.String[] args)
    {
        Connection c = null;

        try {
            // Get the driver registered.
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

            c = DriverManager.getConnection("jdbc:db2:*local");
            DatabaseMetaData dmd = c.getMetaData();

            if (jdbcLevel == 1) {
                System.out.println("No support is provided for getUDTs. Just return.");
                System.exit(1);
            }
        }
    }
}

```

```

ResultSet rs = dmd.getUDTs(null, "CUJOSQL", "SSN%", null);
while (rs.next()) {

    // Fetch all the columns that have been available since the
    // JDBC 2.0 release.
    System.out.println("TYPE_CAT is " + rs.getString("TYPE_CAT"));
    System.out.println("TYPE_SCHEM is " + rs.getString("TYPE_SCHEM"));
    System.out.println("TYPE_NAME is " + rs.getString("TYPE_NAME"));
    System.out.println("CLASS_NAME is " + rs.getString("CLASS_NAME"));
    System.out.println("DATA_TYPE is " + rs.getString("DATA_TYPE"));
    System.out.println("REMARKS is " + rs.getString("REMARKS"));

    // Fetch all the columns that were added in JDBC 3.0.
    if (jdbcLevel > 2) {
        System.out.println("BASE_TYPE is " + rs.getString("BASE_TYPE"));
    }
}
} catch (Exception e) {
    System.out.println("Error: " + e.getMessage());
} finally {
    if (c != null) {
        try {
            c.close();
        } catch (SQLException e) {
            // Ignoring shutdown exception.
        }
    }
}
}
}

```



Exceptions

» The Java^(TM) language uses exceptions to provide error-handling capabilities for its programs. An exception is an event that occurs when you run your program that disrupts the normal flow of instructions.

The Java runtime system and many classes from Java packages throw exceptions in some circumstances by using the throw statement. You can use the same mechanism to throw exceptions in your Java programs.

To find out more about exceptions, see the following sections:

SQLException

The SQLException class and its subtypes provide information about errors and warnings that occur while a data source is being accessed.

SQLWarning

Methods generate an SQLWarning object if the methods cause a database access warning. Methods in the following interfaces can generate SQLWarnings:

- Connection
- Statement and its subtypes, PreparedStatement, and CallableStatement
- ResultSet

DataTruncation

DataTruncation is a subclass of SQLWarning. While SQLWarnings are not thrown, DataTruncation objects are sometimes thrown and attached like other SQLWarning objects.

Silent truncation

The `setMaxFieldSize` statement method allows a maximum field size to be specified for any column. If data truncates because its size has exceeded the maximum field size value, no warning or exception is reported.



SQLException: ➤ The `SQLException` class and its subtypes provide information about errors and warnings that occur while a data source is being accessed.

Unlike most of JDBC, which is defined by interfaces, the exception support is provided in classes. The base class for exceptions that occur while running JDBC applications is `SQLException`. Every method of the JDBC API is declared as being able to throw `SQLException`s. `SQLException` is an extension of `java.lang.Exception` and provides additional information related to failures that happen in a database context. Specifically, the following information is available from an `SQLException`:

- Text description
- `SQLState`
- Error code
- A reference to any other exceptions that also occurred

`ExceptionExample` is a program that properly handles catching an (expected in this case) `SQLException` and dumping all the information that it provides.

Note: JDBC provides a mechanism where exceptions can be chained together. This allows the driver or the database to report multiple errors on a single request. There are currently no instances where the native JDBC driver would do this. This information is only provided as reference and not a clear indication that the driver never does this in the future however.

As noted, `SQLException` objects are thrown when errors occur. This is correct, but is not the complete picture. In practice, the native JDBC driver rarely throws actual `SQLException`s. It throws instances of its own `SQLException` subclasses. This allows you to determine more information about what has actually failed as is shown below.

DB2Exception.java: `DB2Exception` objects are not thrown directly either. This base class is used to hold functionality that is common to all JDBC exceptions. There are two subclasses of this class that are the standard exceptions that JDBC throws. These subclasses are `DB2DBException.java` and `DB2JDBCException.java`. `DB2DBExceptions` are exceptions that are reported to you that have come directly from the database. `DB2JDBCExceptions` are thrown when the JDBC driver finds problems on its own. Splitting the exception class hierarchy in this manner allows you to handle the two types of exceptions differently.

DB2DBException.java: As stated, `DB2DBExceptions` are exceptions that come directly from the database. These are encountered when the JDBC driver makes a call to the CLI and gets back an `SQLERROR` return code. The CLI function `SQLERROR` is called to get the message text, `SQLState`, and vendor code in these cases. The replacement text for the `SQLMessage` is also retrieved and returned to you. The `DatabaseException` class causes an error that the database recognizes and reports to the JDBC driver to build the exception object for.

DB2JDBCException.java: `DB2JDBCExceptions` are generated for error conditions that come from the JDBC driver itself. The functionality of this exception class is fundamentally different; the JDBC driver itself handles message language translation of exception and other issues that the operating system and database handle for exceptions originating within the database. Wherever possible, the JDBC driver adheres to the `SQLStates` of the database. The vendor code for exceptions that the JDBC driver throws is always -99999. `DB2DBExceptions` that are recognized and returned by the CLI layer often also have the -99999 error code. The `JDBCException` class causes an error that the JDBC driver recognizes and builds

the exception for itself. When run during development of the release, the following output was created. Notice that the top of the stack contains DB2JDBCException. This is an indication that the error is being reported from the JDBC driver prior to ever making the request to the database. <<

Example: SQLException: >> This is an example of catching an SQLException and dumping all the information that it provides.

Example: SQLException

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;

public class ExceptionExample {

    public static Connection connection = null;

    public static void main(java.lang.String[] args) {

        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            connection = DriverManager.getConnection("jdbc:db2:*local");

            Statement s = connection.createStatement();
            int count = s.executeUpdate("insert into cujofake.cujofake values(1, 2,3)");

            System.out.println("Did not expect that table to exist.");

        } catch (SQLException e) {
            System.out.println("SQLException exception: ");
            System.out.println("Message:....." + e.getMessage());
            System.out.println("SQLState:...." + e.getSQLState());
            System.out.println("Vendor Code:.." + e.getErrorCode());
            System.out.println("-----");
            e.printStackTrace();
        } catch (Exception ex) {
            System.out.println("An exception other than an SQLException was thrown: ");
            ex.printStackTrace();
        } finally {
            try {
                if (connection != null) {
                    connection.close();
                }
            } catch (SQLException e) {
                System.out.println("Exception caught attempting to shutdown...");
            }
        }
    }
}
```

<<

SQLWarning: >> Methods in the following interfaces generate an SQLWarning object if the methods cause a database access warning:

- Connection
- Statement and its subtypes, PreparedStatement and CallableStatement
- ResultSet

When a method generates an SQLWarning object, the caller is not informed that a data access warning has occurred. The getWarnings method must be called on the appropriate object to retrieve the SQLWarning object. However, the DataTruncation subclass of SQLWarning may be thrown in some

circumstances. It should be noted that the native JDBC driver opts to ignore some database-generated warnings for increased efficiency. For example, a warning is generated by the system when you attempt to retrieve data beyond the end of a `ResultSet` through the `ResultSet.next` method. In this case, the next method is defined to return `false` instead of `true`, informing you of the error. It is unnecessary to create an object to restate this, so the warning is simply ignored.

If multiple data access warnings occur, they are chained to the first one and can be retrieved by calling the `SQLWarning.getNextWarning` method. If there are no more warnings in the chain, `getNextWarning` returns `null`.

Subsequent `SQLWarning` objects continue to be added to the chain until the next statement is processed or, in the case of a `ResultSet` object, when the cursor is repositioned. As a result, all `SQLWarning` objects in the chain are removed.

Using `Connection`, `Statement`, and `ResultSet` objects can cause `SQLWarnings` to be generated. `SQLWarnings` are informational messages indicating that while a particular operation has completed successfully, there might be other information of which you should be aware. `SQLWarnings` are an extension of the `SQLException` class, but they are not thrown. They are instead attached to the object that causes their generation. When an `SQLWarning` is generated, nothing happens to inform the application that the warning has been generated. Your application must actively request warning information.

Like `SQLExceptions`, `SQLWarnings` can be chained to one another. You can call the `clearWarnings` method on a `Connection`, `Statement`, or `ResultSet` object to clear the warnings for that object.

Note: Calling the `clearWarnings` method does not clear all warnings. It only clears the warnings that are associated with a particular object.

The JDBC driver clears `SQLWarning` objects at specific times if you do not clear them manually. `SQLWarning` objects are cleared when the following actions are taken:

- For the `Connection` interface, warnings are cleared on the creation of a new `Statement`, `PreparedStatement`, or `CallableStatement` object.
- For the `Statement` interface, warnings are cleared when the next statement is processed (or when the statement is processed again for `PreparedStatements` and `CallableStatements`).
- For the `ResultSet` interface, warnings are cleared when the cursor is repositioned. <<

DataTruncation: >> `DataTruncation` is a subclass of `SQLWarning`. While `SQLWarnings` are not thrown, `DataTruncation` objects are sometimes thrown and attached like other `SQLWarning` objects. `DataTruncation` objects provide additional information beyond what is returned by an `SQLWarning`. The available information includes the following:

- The number of bytes of data that should have been transferred.
- The column or parameter index that was truncated.
- Whether the index is for a parameter or a `ResultSet` column.
- Whether the truncation happened when reading from the database or writing to it.
- The amount of data that was actually transferred.

In some instances, the information can be deciphered, but situations arise that are not completely intuitive. For example, if the `PreparedStatement`'s `setFloat` method is used to insert a value into a column that holds integer values, a `DataTruncation` may result because the float may be larger than the largest value that the column can hold. In these situations, the byte counts for truncation do not make sense, but it is important for the driver to provide the truncation information.

Report set() and update() methods: There is a subtle difference between JDBC drivers. Some drivers such as the native and Toolbox JDBC drivers catch and report data truncation issues at the time of the

parameter setting. This is done either on the PreparedStatement set method or the ResultSet update method. Other drivers report the problem at the time of processing the statement and is accomplished by the execute, executeQuery, or updateRow methods.

Failing to report the problem at the time that you provide incorrect data instead of at the time that processing cannot continue any further offers a couple of advantages:

- The failure can be addressed in your application when you have a problem instead of addressing the problem at processing time.
- By checking when setting the parameters, the JDBC driver can ensure that the values that are handed to the database at statement processing time are valid. This allows the database to optimize its work and processing can be completed faster.

ResultSet.update() methods throw DataTruncation exceptions: In some past releases, ResultSet.update() methods posted warnings when truncation conditions existed. This case occurs when the data value is going to be inserted into the database. The specification dictates that JDBC drivers should throw exceptions in these cases. As a result, the JDBC driver works in this manner.

There should be no significant difference between handling a ResultSet update function that receives a data truncation error and handling a prepared statement parameter set for an update or insert statement that receives an error. In both cases, the problem is identical; you provided data that does not fit where you wanted it.

NUMERIC and DECIMAL truncate to the right side of a decimal point silently. This is how both JDBC for UDB NT works and how interactive SQL on an iSeries server works.

Note: No value is rounded when a data truncation occurs. Any fractional portion of a parameter that does not fit in a NUMERIC or DECIMAL column is simply lost without warning.

The following are examples, assuming that the value in the values clause is actually a parameter being set on a prepared statement:

```
create table cujosql.test (col1 numeric(4,2))
a) insert into cujosql.test values(22.22) // works - inserts 22.22
b) insert into cujosql.test values(22.223) // works - inserts 22.22
c) insert into cujosql.test values(22.227) // works - inserts 22.22
d) insert into cujosql.test values(322.22) // fails - Conversion error on assignment to column COL1.
```

Difference between a data truncation warning and a data truncation exception

The specification states that data truncation on a value to be written to the database should throw an exception. If data truncation is not performed on the value being written to the database, a warning should be generated. This means that the point at which a data truncation situation is identified, you must also be aware of the statement type that the data truncation is processing. Given this as a requirement, the following lists the behavior of several SQL statement types:

- In a SELECT statement, query parameters never damage database content. Therefore, data truncation situations are always handled by posting warnings.
- In VALUES INTO and SET statements, the input values are only used to generate output values. As a result, warnings are issued.
- In a CALL statement, the JDBC driver cannot determine what a stored procedure does with a parameter. Exceptions are always thrown when a stored procedure parameter truncates.
- All other statement types throw exceptions rather than post warnings.

Data truncation property for Connection and DataSource: There has been a data truncation property available for many releases. The default for that property is true, meaning that data truncation issues are checked and warnings are posted or exceptions are thrown. The property is provided for convenience and

performance in cases where you are not concerned that a value does not fit into the database column. You want the driver to put as much of the value as it can into the column.

Data truncation property only affects character and binary-based data types: A couple releases ago, the data truncation property determined whether data truncation exceptions could be thrown. The data truncation property was put in place to have JDBC applications not worry about a value getting truncated when the truncation was not important to them. There are few cases where you would want either the value 00 or 10 stored in the database when applications attempted to insert 100 into a DECIMAL(2,0). Therefore, the JDBC driver's data truncation property was changed to only honor situations where the parameter is for character-based types such as CHAR, VARCHAR, CHAR FOR BIT DATA, and VARCHAR FOR BIT DATA.

Data truncation property is only applied to parameters: The data truncation property is a setting of the JDBC driver and not of the database. As a result, it has no effect on statement literals. For example, the following statements that are processed to insert a value into a CHAR(8) column in the database still fail with the data truncation flag set to false (assume that connection is a java.sql.Connection object allocated elsewhere).

```
Statement stmt = connection.createStatement();
Stmt.executeUpdate("create table cujosql.test (col1 char(8))");
Stmt.executeUpdate("insert into cujosql.test values('dettinger')");
// Fails as the value does not fit into database column.
```

Native JDBC driver throws exceptions for insignificant data truncation: The native JDBC driver does not look at the data that you provide for parameters. Doing so only slows down processing. However, there can be situations where it does not matter to you that a value truncates, but you have not set the data truncation connection property to false.

For example, 'dettinger ', a char(10) that is passed, throws an exception even though everything important about the value fits. This does happen to be how JDBC for UDB NT works; however, it is not the behavior you would get if you passed the value as a literal in an SQL statement. In this case, the database engine would throw out the additional spaces quietly.

The problems with the JDBC driver not throwing an exception are the following:

- Performance overhead is extensive on every set method, whether needed or not. For the majority of cases where there would be no benefit, there is considerable performance overhead on a function as common as setString().
- Your workaround is trivial, for example, calling the trim function on the string value passed in.
- There are issues with the database column to take into account. A space in CCSID 37 is not at all a space in CCSID 65535, or 13488.

Silent truncation: The setMaxFieldSize statement method allows a maximum field size to be specified for any column. If data truncates because its size has exceeded the maximum field size value, no warning or exception is reported. This method, like the data truncation property previously mentioned, only affects character-based types such as CHAR, VARCHAR, CHAR FOR BIT DATA, and VARCHAR FOR BIT DATA.



Transactions

➤ A **transaction** is a logical unit of work. To complete a logical unit of work, several actions may have to be taken against a database. Transactional support allows applications to ensure the following:

- All the steps to complete a logical unit of work are followed.
- When one of the steps to the unit of work files fails, all the work done as part of that logical unit of work can be undone and the database can return to its previous state before the transaction began.

Transactions are used to provide data integrity, correct application semantics, and a consistent view of data during concurrent access. All Java^(TM) Database Connectivity (JDBC) compliant drivers must support transactions.

Note: This section only discusses local transactions and the standard JDBC concept of transactions. Java and the native JDBC driver support the Java Transaction API (JTA), distributed transactions, and the two-phase commit protocol (2PC).

All transactional work is handled at the Connection object level. When the work for a transaction completes, it can be finalized by calling the commit method. If the application aborts the transaction, the rollback method is called.

All Statement objects under a connection are a part of the transaction. This means is that if an application creates three Statement objects and uses each object to make changes to the database, when a commit or rollback call happens, the work for all three statements either becomes permanent or is discarded.

The commit and rollback SQL statements are used to finalize transactions when working purely with SQL. These SQL statements cannot be dynamically prepared and you should not attempt to use them in your JDBC applications to complete transactions.

To use transactions in your application, see the following:

Auto-commit mode

JDBC uses an auto-commit mode where every update to the database is immediately made permanent.

Transaction isolation levels


Transaction isolation levels specify what data is visible to statements within a transaction and directly impact the level of concurrent access.

Savepoints

Savepoints are checkpoints that the application can roll back to without throwing away the entire transaction. Find the following information about Savepoints:

- Set and roll back to savepoints
- Release a savepoint



Auto-commit mode:  By default, JDBC uses an operation mode called auto-commit. This means that every update to the database is immediately made permanent. Any situation where a logical unit of work requires more than one update to the database cannot be done safely in auto-commit mode. If something happens to the application or the system after one update is made and before any other updates are made, the first change cannot be undone when running in auto-commit mode.

Since changes are instantly made permanent in auto-commit mode, there is no need for the application to call the commit method or the rollback method. This makes applications easier to write.

Auto-commit mode can be enabled and disabled dynamically during a connection's existence. Auto-commit is enabled in the following way, assuming that data source already exists:

```
Connection connection = dataSource.getConnection();  
  
Connection.setAutoCommit(false);           // Disables auto-commit.
```

If the auto-commit setting is changed in the middle of a transaction, any pending work is automatically committed. An SQLException is generated if auto-commit is enabled for a connection that is part of a distributed transaction. <<

Transaction isolation levels: >> Transaction isolation levels specify what data is visible to statements within a transaction. These levels directly impact the level of concurrent access by defining what interaction is possible between transactions against the same target data source.

Database anomalies: Database anomalies are generated results that seem incorrect when looked at from the scope of a single transaction, but are correct when looked at from the scope of all transactions. The different types of database anomalies are described as follows:

- **Dirty** reads occur when:
 1. Transaction A inserts a row into a table.
 2. Transaction B reads the new row.
 3. Transaction A rolls back.

Transaction B may have done work to the system based on the row inserted by transaction A, but that row never became a permanent part of the database.

- **Nonrepeatable** reads occur when:
 1. Transaction A reads a row.
 2. Transaction B changes the row.
 3. Transaction A reads the same row a second time and gets the new results.
- **Phantom** reads occur when:
 1. Transaction A reads all rows that satisfy a WHERE clause on an SQL query.
 2. Transaction B inserts an additional row that satisfies the WHERE clause.
 3. Transaction A re-evaluates the WHERE condition and picks up the additional row.

Note: DB2/400 does not always expose the application to the allowable database anomalies at the prescribed levels due to its locking strategies.

JDBC transaction isolation levels: There are five levels of transaction isolation in the IBM Developer Kit for Java JDBC API. Listed from least to most restrictive, they are as follows:

JDBC_TRANSACTION_NONE

This is a special constant indicating that the JDBC driver does not support transactions.

JDBC_TRANSACTION_READ_UNCOMMITTED

This level allows transactions to see uncommitted changes to the data. All database anomalies are possible at this level.

JDBC_TRANSACTION_READ_COMMITTED

This level means that any changes made inside a transaction are not visible outside it until the transaction is committed. This prevents dirty reads from being possible.

JDBC_TRANSACTION_REPEATABLE_READ

This level means that rows that are read retain locks so that another transaction cannot change them when the transaction is not completed. This disallows dirty reads and nonrepeatable reads. Phantom read are still possible.

JDBC_TRANSACTION_SERIALIZABLE

Tables are locked for the transaction so that WHERE conditions cannot be changed by other transactions that add values to or remove values from a table. This prevents all types of database anomalies.

The setTransactionIsolation method can be used to change the transaction isolation level for a connection.

Considerations: A common misinterpretation is that the JDBC specification defines the five transactional levels previously mentioned. It is commonly thought that the TRANSACTION_NONE value represents the concept of running without commitment control. The JDBC specification does not define TRANSACTION_NONE in the same manner. TRANSACTION_NONE is defined in the JDBC specification as a level where the driver does not support transactions and is not a JDBC-compliant driver. The NONE level is never reported when the getTransactionIsolation method is called.

The issue is marginally complicated by the fact that a JDBC driver's default transaction isolation level is defined by the implementation. The default level of transaction isolation for the native JDBC driver default transaction isolation level is NONE. This allows the driver to work with files that do not have journals and you are not required to make any specifications such as files in the QGPL library.

The native JDBC driver allows you to pass JDBC_TRANSACTION_NONE to the setTransactionIsolation method or specify none as a connection property. However, the getTransactionIsolation method always reports JDBC_TRANSACTION_READ_UNCOMMITTED when the value is none. It is your application's responsibility to keep track of what level you are running if it is a requirement in your application.

In past releases, the JDBC driver would handle your specifying true for auto-commit by changing the transaction isolation level to none because the system did not have a concept of a true auto-commit mode. This was a close approximation of the functionality, but did not provide the correct results for all scenarios. This is not done anymore; the database decouples the concept of auto-commit from the concept of a transaction isolation level. Therefore, it is completely valid to run at the JDBC_TRANSACTION_SERIALIZABLE level with auto-commit being enabled. The only scenario that is not valid is to run at the JDBC_TRANSACTION_NONE level and not be in auto-commit mode. Your application cannot take control over commit boundaries when the system is not running with a transaction isolation level.

Transaction isolation levels between the JDBC specification and the iSeries platform: The iSeries platform has common names for its transaction isolation levels that do not match those names provided by the JDBC specification. The following table matches the names used by the iSeries platform, but are not equivalent to those used by the JDBC specification:

JDBC level*	iSeries level
JDBC_TRANSACTION_NONE	*NONE or *NC
JDBC_TRANSACTION_READ_UNCOMMITTED	*CHG or *UR
JDBC_TRANSACTION_READ_COMMITTED	*CS
JDBC_TRANSACTION_REPEATABLE_READ	*ALL or *RS
JDBC_TRANSACTION_SERIALIZABLE	*RR

* In this table, the JDBC_TRANSACTION_NONE value is lined up with the iSeries levels *NONE and *NC for clarity. This is not a direct specification-to-iSeries level match. <<

Savepoints: >> Savepoints allow the setting of "staging points" in a transaction. Savepoints are checkpoints that the application can roll back to without throwing away the entire transaction. Savepoints are new in JDBC 3.0, meaning that the application must run on Java^(TM) Development Kit (JDK) 1.4 to use them. Moreover, savepoints are new to the Developer Kit for Java, meaning that savepoints are not supported if JDK 1.4 is not used with previous releases of the Developer Kit for Java.

Note: The system provides SQL statements for working with savepoints. It is advised that JDBC applications do not use these statements directly in an application. Doing so may work, but the JDBC driver loses its ability to track the your savepoints when this is done. At a minimum, mixing the two models (that is, using your own savepoint SQL statements and using the JDBC API) should be avoided.


Set and roll back to savepoints: Savepoints can be set throughout the work of a transaction. The application can then roll back to any of these savepoints if something goes wrong and continue processing from that point. In the following example, the application inserts the value FIRST into a database table. After that, a savepoint is set and another value, SECOND, is inserted into the database. A rollback to the savepoint is issued and undoes the work of inserting SECOND, but leaves FIRST as part of the pending transaction. Finally, the value THIRD is inserted and the transaction is committed. The database table contains the values FIRST and THIRD.

Example: Set and roll back to savepoints


Note: Read the Code example disclaimer for important legal information.


```
Statement s = Connection.createStatement();
s.executeUpdate("insert into table1 values ('FIRST')");
Savepoint pt1 = connection.setSavepoint("FIRST SAVEPOINT");
s.executeUpdate("insert into table1 values ('SECOND')");
connection.rollback(pt1); // Undoes most recent insert.
s.executeUpdate("insert into table1 values ('THIRD')");
connection.commit();
```

Although it is unlikely to cause problems to set savepoints while in auto-commit mode, they cannot be rolled back as their lives end at the end of a transaction.

Release a savepoint: Savepoints can be released by the application with the `releaseSavepoint` method on the `Connection` object. Once a savepoint has been released, attempting to roll back to it results in an exception. When a transaction commits or rolls back, all savepoints automatically release. When a savepoint is rolled back, other savepoints that follow it are also released. 

Distributed transactions

 Typically, transactions in Java[™] Database Connectivity (JDBC) are local. This means that a single connection performs all the work of the transaction and that the connection can only work on one transaction at a time. When all the work for that transaction has been completed or has failed, `commit` or `rollback` is called to make the work permanent, and a new transaction can begin.

There is advanced support for transactions available in Java that provides functionality beyond local transactions. This support is fully specified by the Java Transaction API (JTA) 1.0.1 specification .

The Java Transaction API (JTA) has support for complex transactions. It also provides support for decoupling transactions from `Connection` objects. As JDBC is modeled after the Object Database Connectivity (ODBC) and the X/Open Call Level Interface (CLI) specifications, JTA is modeled after the X/Open Extended Architecture (XA) specification. JTA and JDBC work together to decouple transactions from `Connection` objects. By decoupling transactions from `Connection` objects, this allows you to have a single connection work on multiple transactions concurrently. Conversely, it allows you to have multiple `Connections` work on a single transaction.

Note: If you are planning to work with JTA, refer to *Get started with JDBC* for more information about required Java Archive (JAR) files in your extensions classpath. You want both the JDBC 2.0 optional package and the JTA JAR files (these files are found automatically by the JDK if you are running JDK 1.4). These are not found by default.

Transactions with JTA: When JTA and JDBC are used together, there are a series of steps between them to accomplish transactional work. Support for XA is provided through the `XADataSource` class. This class contains support for setting up connection pooling exactly the same way as its `ConnectionPoolDataSource` superclass.

With an XADataSource instance, you can retrieve an XAConnection object. The XAConnection object serves as a container for both the JDBC Connection object and an XAResource object. The XAResource object is designed to handle XA transactional support. XAResource handles transactions through objects called transaction IDs (XIDs).

The XID is an interface that you must implement. It represents a Java mapping of the XID structure of the X/Open transaction identifier. This object contains three parts:

- A global transaction's format ID
- A global transaction ID
- A branch qualifier

See the JTA specification for complete details on this interface.

Example: Use JTA to handle a transaction shows how to use JTA to handle a transaction in an application.

Use UDBXDataSource support for pooling and distributed transactions: The Java Transaction API support provides direct support for connection pooling. UDBXDataSource is an extension of a ConnectionPoolDataSource, allowing application access to pooled XAConnection objects. Since UDBXDataSource is a ConnectionPoolDataSource, the configuration and use of the UDBXDataSource is the same as that described in Use DataSource support for object pooling.

XADataSource properties: In addition to the properties provided by the ConnectionPoolDataSource, the XADataSource interface provides the following properties:

Set method (data type)	Values	Description
setLockTimeout (int)	0 or any positive value	Any positive value is a valid lock timeout (in seconds) at the transaction level. A lock timeout of 0 means that there is no lock timeout value enforced at the transaction level, although there may be one enforced at other levels (the job or the table). The default value is 0.
setTransactionTimeout (int)	0 or any positive value	Any positive value is a valid transaction timeout (in seconds). A transaction timeout of 0 means that there is no transaction timeout value enforced. If the transaction is active for longer than the timeout value, it is marked rollback only, and subsequent attempts to perform work under it causes an exception to occur. The default value is 0.

ResultSets and transactions: Besides demarcating the start and end of a transaction as shown in the previous example, transactions can be suspended for a time and resumed later. This provides a number of scenarios for ResultSet resources that are created during a transaction.

Simple transaction end: When you end a transaction, all open ResultSets that were created under that transaction automatically close. It is recommended that you explicitly close your ResultSets when you are

finished using them to ensure maximum parallel processing. However, an exception results if any ResultSets that were opened during a transaction are accessed after the XAResource.end call is made.

See Example: End a transaction that shows this behavior.

Suspend and resume: While a transaction is suspended, access to a ResultSet created while the transaction was active is not allowed and results in an exception. However, once the transaction is resumed, the ResultSet is available again and remains in the same state it was in before the transaction was suspended.

See Example: Suspend and resume a transaction that shows this behavior.

Effecting suspended ResultSets: While a transaction is suspended, the ResultSet cannot be accessed. However, Statement objects can be reprocessed under another transaction to perform work. Because JDBC Statement objects can have only one ResultSet at a time (excluding the JDBC 3.0 support for multiple concurrent ResultSets from a stored procedure call), the ResultSet for the suspended transaction must be closed to fulfill the request of the new transaction. This is exactly what happens.

See Example: Suspended ResultSets that shows this behavior.

Note: Although JDBC 3.0 allows a Statement to have multiple ResultSets open simultaneously for a stored procedure call, they are treated as a single unit and all of them close if the Statement is reprocessed under a new transaction. Currently, it is not possible to have ResultSets from two transactions active simultaneously for a single statement.

Multiplexing: The JTA API is designed to decouple transactions from JDBC connections. This API allows you to have either multiple connections work on a single transaction or a single connection work on multiple transactions concurrently. This is called **multiplexing** and many complex tasks can be performed that cannot be accomplished with JDBC alone.

This example shows multiple connections working on a single transaction.

This example shows a single connection with multiple transactions taking place at once.

For further information on using JTA, see the JTA specification. The JDBC 3.0 specification also contains information on how these two technologies work together to support distributed transactions.

Two-phase commit and transaction logging: The JTA APIs externalize the responsibilities of the distributed two-phase commit protocol completely to the application. As the examples have shown, when using JTA and JDBC to access a database under a JTA transaction, the application uses the XAResource.prepare() and XAResource.commit() methods or just the XAResource.commit() method to commit the changes.

In addition, when accessing multiple distinct databases using a single transaction, it is the application's responsibility to ensure that the two-phase commit protocol and any associated logging required for transaction atomicity across those databases are performed. Typically, the two-phase commit processing across multiple databases (that is, XAResources) and its logging are performed under the control of an application server or transaction monitor so that the application itself does not actually concern itself with these issues.

For example, the application may call some commit() method or return from its processing with no errors. The underlying application server or transaction monitor would then begin processing for each database (XAResource) that participated in the single distributed transaction.

The application server would use extensive logging during the two-phase commit processing. It would call the `XAResource.prepare()` method in turn for each participant database (`XAResource`), followed by a call to the `XAResource.commit()` method for each participant database (`XAResource`).

If a failure occurs during this processing, the application server's transaction monitor logs allow the application server itself to subsequently use the JTA APIs to recover the distributed transaction. This recovery, under the control of the application server or transaction monitor, allows the application server to get the transaction to a known state at each participant database (`XAResource`). This ensures a well-known state of the entire distributed transaction across all participant databases. <<

Example: Use JTA to handle a transaction: >> This is an example of how to use the Java[™] Transaction API (JTA) to handle a transaction in an application.

Example: Use JTA to handle a transaction

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;

public class JTACommit {

    public static void main(java.lang.String[] args) {
        JTACommit test = new JTACommit();

        test.setup();
        test.run();
    }

    /**
     * Handle the previous cleanup run so that this test can recommence.
     */
    public void setup() {

        Connection c = null;
        Statement s = null;
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            c = DriverManager.getConnection("jdbc:db2:*local");
            s = c.createStatement();

            try {
                s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
            } catch (SQLException e) {
                // Ignore... does not exist
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR (50))");
            s.close();
        } finally {
            if (c != null) {
```

```

        c.close();
    }
}

/**
 * This test uses JTA support to handle transactions.
 */
public void run() {
    Connection c = null;

    try {
        Context ctx = new InitialContext();

        // Assume the data source is backed by a UDBXADatasource.
        UDBXADatasource ds = (UDBXADatasource) ctx.lookup("XADatasource");

        // From the DataSource, obtain an XAConnection object that
        // contains an XAResource and a Connection object.
        XAConnection xaConn = ds.getXAConnection();
        XAResource xaRes = xaConn.getXAResource();
        Connection c = xaConn.getConnection();

        // For XA transactions, a transaction identifier is required.
        // An implementation of the XID interface is not included with the
        // JDBC driver. See Transactions with JTA for a description of
        // this interface to build a class for it.
        Xid xid = new XidImpl();

        // The connection from the XAResource can be used as any other
        // JDBC connection.
        Statement stmt = c.createStatement();

        // The XA resource must be notified before starting any
        // transactional work.
        xaRes.start(xid, XAResource.TMNOFLAGS);

        // Standard JDBC work is performed.
        int count = stmt.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('JTA is pretty fun.')");

        // When the transaction work has completed, the XA resource must
        // again be notified.
        xaRes.end(xid, XAResource.TMSUCCESS);

        // The transaction represented by the transaction ID is prepared
        // to be committed.
        int rc = xaRes.prepare(xid);

        // The transaction is committed through the XAResource.
        // The JDBC Connection object is not used to commit
        // the transaction when using JTA.
        xaRes.commit(xid, false);

    } catch (Exception e) {

```

```

        System.out.println("Something has gone wrong.");
        e.printStackTrace();
    } finally {
        try {
            if (c != null)
                c.close();
        } catch (SQLException e) {
            System.out.println("Note: Cleanup exception.");
            e.printStackTrace();
        }
    }
}
}
}

```



Example: Multiple connections that work on a transaction: ➤ This is an example of how to use multiple connections working on a single transaction.

Example: Multiple connections that work on a transaction

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;
public class JTAMultiConn {
    public static void main(java.lang.String[] args) {
        JTAMultiConn test = new JTAMultiConn();
        test.setup();
        test.run();
    }
}
/**
 * Handle the previous cleanup run so that this test can recommence.
 */
public void setup() {
    Connection c = null;
    Statement s = null;
    try {
        Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        c = DriverManager.getConnection("jdbc:db2:*local");
        s = c.createStatement();
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
        }
        catch (SQLException e) {
            // Ignore... does not exist
        }
        s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR
            (50))");
        s.close();
    }
    finally {
        if (c != null) {
            c.close();
        }
    }
}
}
/**
 * This test uses JTA support to handle transactions.

```

```

*/
public void run() {
    Connection c1 = null;
    Connection c2 = null;
    Connection c3 = null;
    try {
        Context ctx = new InitialContext();
        // Assume the data source is backed by a UDBXADDataSource.
        UDBXADDataSource ds = (UDBXADDataSource)
            ctx.lookup("XADataSource");
        // From the DataSource, obtain some XAConnection objects that
        // contain an XAResource and a Connection object.
        XAConnection xaConn1 = ds.getXAConnection();
        XAConnection xaConn2 = ds.getXAConnection();
        XAConnection xaConn3 = ds.getXAConnection();
        XAResource xaRes1 = xaConn1.getXAResource();
        XAResource xaRes2 = xaConn2.getXAResource();
        XAResource xaRes3 = xaConn3.getXAResource();
        c1 = xaConn1.getConnection();
        c2 = xaConn2.getConnection();
        c3 = xaConn3.getConnection();
        Statement stmt1 = c1.createStatement();
        Statement stmt2 = c2.createStatement();
        Statement stmt3 = c3.createStatement();
        // For XA transactions, a transaction identifier is required.
        // Support for creating XIDs is again left to the application
        // program.
        Xid xid = JDXATest.xidFactory();
        // Perform some transactional work under each of the three
        // connections that have been created.
        xaRes1.start(xid, XAResource.TMNOFLAGS);
        int count1 = stmt1.executeUpdate("INSERT INTO " + tableName + "VALUES('Value 1-A')");
        xaRes1.end(xid, XAResource.TMNOFLAGS);

        xaRes2.start(xid, XAResource.TMJOIN);
        int count2 = stmt2.executeUpdate("INSERT INTO " + tableName + "VALUES('Value 1-B')");
        xaRes2.end(xid, XAResource.TMNOFLAGS);

        xaRes3.start(xid, XAResource.TMJOIN);
        int count3 = stmt3.executeUpdate("INSERT INTO " + tableName + "VALUES('Value 1-C')");
        xaRes3.end(xid, XAResource.TMSUCCESS);
        // When completed, commit the transaction as a single unit.
        // A prepare() and commit() or 1 phase commit() is required for
        // each separate database (XAResource) that participated in the
        // transaction. Since the resources accessed (xaRes1, xaRes2, and xaRes3)
        // all refer to the same database, only one prepare or commit is required.
        int rc = xaRes.prepare(xid);
        xaRes.commit(xid, false);
    }
    catch (Exception e) {
        System.out.println("Something has gone wrong.");
        e.printStackTrace();
    }
    finally {
        try {
            try {
                if (c1 != null) {
                    c1.close();
                }
            }
            catch (SQLException e) {
                System.out.println("Note: Cleanup exception " +
                    e.getMessage());
            }
        }
        try {
            if (c2 != null) {
                c2.close();
            }
        }
    }
}

```

```

    }
    catch (SQLException e) {
        System.out.println("Note: Cleanup exception " +
            e.getMessage());
    }
    try {
        if (c3 != null) {
            c3.close();
        }
    }
    catch (SQLException e) {
        System.out.println("Note: Cleanup exception " +
            e.getMessage());
    }
}
}
}
}

```



Example: Use a connection with multiple transactions: ➤ This is an example of how to use a single connection with multiple transactions.

Example: Use a connection with multiple transactions

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;

public class JTAMultiTx {

    public static void main(java.lang.String[] args) {
        JTAMultiTx test = new JTAMultiTx();

        test.setup();
        test.run();
    }

    /**
     * Handle the previous cleanup run so that this test can recommence.
     */
    public void setup() {

        Connection c = null;
        Statement s = null;
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            c = DriverManager.getConnection("jdbc:db2:*local");
            s = c.createStatement();

            try {
                s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
            } catch (SQLException e) {
                // Ignore... does not exist
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR (50))");

            s.close();
        } finally {

```

```

        if (c != null) {
            c.close();
        }
    }
}

/**
 * This test uses JTA support to handle transactions.
 */
public void run() {
    Connection c = null;

    try {
        Context ctx = new InitialContext();

        // Assume the data source is backed by a UDBXADDataSource.
        UDBXADDataSource ds = (UDBXADDataSource) ctx.lookup("XADDataSource");

        // From the DataSource, obtain an XAConnection object that
        // contains an XAResource and a Connection object.
        XAConnection xaConn = ds.getXAConnection();
        XAResource xaRes = xaConn.getXAResource();
        Connection c = xaConn.getConnection();
        Statement stmt = c.createStatement();

        // For XA transactions, a transaction identifier is required.
        // This is not meant to imply that all the XIDs are the same.
        // Each XID must be unique to distinguish the various transactions
        // that occur.
        // Support for creating XIDs is again left to the application
        // program.
        Xid xid1 = JDXTTest.xidFactory();
        Xid xid2 = JDXTTest.xidFactory();
        Xid xid3 = JDXTTest.xidFactory();

        // Do work under three transactions for this connection.
        xaRes.start(xid1, XAResource.TMNOFLAGS);
        int count1 = stmt.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Value 1-A')");
        xaRes.end(xid1, XAResource.TMNOFLAGS);

        xaRes.start(xid2, XAResource.TMNOFLAGS);
        int count2 = stmt.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Value 1-B')");
        xaRes.end(xid2, XAResource.TMNOFLAGS);

        xaRes.start(xid3, XAResource.TMNOFLAGS);
        int count3 = stmt.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Value 1-C')");
        xaRes.end(xid3, XAResource.TMNOFLAGS);

        // Prepare all the transactions
        int rc1 = xaRes.prepare(xid1);
        int rc2 = xaRes.prepare(xid2);
        int rc3 = xaRes.prepare(xid3);

        // Two of the transactions commit and one rolls back.
        // The attempt to insert the second value into the table is
        // not committed.
        xaRes.commit(xid1, false);
        xaRes.rollback(xid2);
        xaRes.commit(xid3, false);

    } catch (Exception e) {
        System.out.println("Something has gone wrong.");
        e.printStackTrace();
    } finally {
        try {

```



```

        if (c != null)
            c.close();
    } catch (SQLException e) {
        System.out.println("Note: Cleanup exception.");
        e.printStackTrace();
    }
}
}
}

```



Example: Suspended ResultSets: ➤ This is an example of the how a Statement object is reprocessed under another transaction to perform work.

Example: Suspended ResultSets

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;

public class JTATxEffct {

    public static void main(java.lang.String[] args) {
        JTATxEffct test = new JTATxEffct();

        test.setup();
        test.run();
    }

    /**
     * Handle the previous cleanup run so that this test can recommence.
     */
    public void setup() {

        Connection c = null;
        Statement s = null;
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            c = DriverManager.getConnection("jdbc:db2:*local");
            s = c.createStatement();

            try {
                s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
            } catch (SQLException e) {
                // Ignore... does not exist
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR (50))");
            s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Fun with JTA')");
            s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('JTA is fun.')");
        }
    }
}

```

```

        s.close();
    } finally {
        if (c != null) {
            c.close();
        }
    }
}

/**
 * This test uses JTA support to handle transactions.
 */
public void run() {
    Connection c = null;

    try {
        Context ctx = new InitialContext();

        // Assume the data source is backed by a UDBXADDataSource.
        UDBXADDataSource ds = (UDBXADDataSource) ctx.lookup("XADataSource");

        // From the DataSource, obtain an XAConnection object that
        // contains an XAResource and a Connection object.
        XAConnection xaConn = ds.getXAConnection();
        XAResource xaRes = xaConn.getXAResource();
        Connection c = xaConn.getConnection();

        // For XA transactions, a transaction identifier is required.
        // An implementation of the XID interface is not included with
        // the JDBC driver. See Transactions with JTA
        // for a description of this interface to build a
        // class for it.
        Xid xid = new XidImpl();

        // The connection from the XAResource can be used as any other
        // JDBC connection.
        Statement stmt = c.createStatement();

        // The XA resource must be notified before starting any
        // transactional work.
        xaRes.start(xid, XAResource.TMNOFLAGS);

        // Create a ResultSet during JDBC processing and fetch a row.
        ResultSet rs = stmt.executeUpdate("SELECT * FROM CUJOSQL.JTATABLE");
        rs.next();

        // The end method is called with the suspend option. The
        // ResultSets associated with the current transaction are 'on hold'.
        // They are neither gone nor accessible in this state.
        xaRes.end(xid, XAResource.TMSUSPEND);

        // In the meantime, other work can be done outside the transaction.
        // The ResultSets under the transaction can be closed if the
        // Statement object used to create them is reused.
    }
}

```

```

ResultSet nonXARS = stmt.executeQuery("SELECT * FROM CUJOSQL.JTATABLE");
while (nonXARS.next()) {
    // Process here...
}

// Attempt to go back to the suspended transaction. The suspended
// transaction's ResultSet has disappeared because the statement
// has been processed again.
xaRes.start(newXid, XAResource.TMRESUME);
try {
    rs.next();
} catch (SQLException ex) {
    System.out.println("This exception is expected. The ResultSet closed due to another process.")
}

// When the transaction had completed, end it
// and commit any work under it.
xaRes.end(xid, XAResource.TMNOFLAGS);
int rc = xaRes.prepare(xid);
xaRes.commit(xid, false);

} catch (Exception e) {
    System.out.println("Something has gone wrong.");
    e.printStackTrace();
} finally {
    try {
        if (c != null)
            c.close();
    } catch (SQLException e) {
        System.out.println("Note: Cleanup exception.");
        e.printStackTrace();
    }
}
}
}
}

```



Example: End a transaction: ➤ This is an example of ending a transaction in your application.

Example: End a transaction

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;

public class JTATxEnd {

```

```

public static void main(java.lang.String[] args) {
    JTATxEnd test = new JTATxEnd();

    test.setup();
    test.run();
}

/**
 * Handle the previous cleanup run so that this test can recommence.
 */
public void setup() {

    Connection c = null;
    Statement s = null;
    try {
        Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        c = DriverManager.getConnection("jdbc:db2:*local");
        s = c.createStatement();

        try {
            s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
        } catch (SQLException e) {
            // Ignore... does not exist
        }

        s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR (50))");
        s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Fun with JTA')");
        s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('JTA is fun.')");

        s.close();
    } finally {
        if (c != null) {
            c.close();
        }
    }
}

/**
 * This test use JTA support to handle transactions.
 */
public void run() {
    Connection c = null;

    try {
        Context ctx = new InitialContext();

        // Assume the data source is backed by a UDBXADatasource.
        UDBXADatasource ds = (UDBXADatasource) ctx.lookup("XADataSource");

        // From the DataSource, obtain an XAConnection object that
        // contains an XAResource and a Connection object.
        XAConnection xaConn = ds.getXAConnection();
        XAResource xaRes = xaConn.getXAResource();
        Connection c = xaConn.getConnection();
    }
}

```

```

// For XA transactions, transaction identifier is required.
// An implementation of the XID interface is not included
// with the JDBC driver. See Transactions with JTA for a
// description of this interface to build a class for it.
Xid xid = new XidImpl();

// The connection from the XAResource can be used as any other
// JDBC connection.
Statement stmt = c.createStatement();

// The XA resource must be notified before starting any
// transactional work.
xaRes.start(xid, XAResource.TMNOFLAGS);

// Create a ResultSet during JDBC processing and fetch a row.
ResultSet rs = stmt.executeUpdate("SELECT * FROM CUJOSQL.JTATABLE");
rs.next();

// When the end method is called, all ResultSet cursors close.
// Accessing the ResultSet after this point results in an
// exception being thrown.
xaRes.end(xid, XAResource.TMNOFLAGS);

try {
    String value = rs.getString(1);
    System.out.println("Something failed if you receive this message.");
} catch (SQLException e) {
    System.out.println("The expected exception was thrown.");
}

// Commit the transaction to ensure that all locks are
// released.
int rc = xaRes.prepare(xid);
xaRes.commit(xid, false);

} catch (Exception e) {
    System.out.println("Something has gone wrong.");
    e.printStackTrace();
} finally {
    try {
        if (c != null)
            c.close();
    } catch (SQLException e) {
        System.out.println("Note: Cleanup exception.");
        e.printStackTrace();
    }
}
}
}
}
}

```



Example: Suspend and resume a transaction: ➤ This is an example of a transaction that is suspended and then is resumed.

Example: Suspend and resume a transaction

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import javax.sql.*;
import java.util.*;
import javax.transaction.*;
import javax.transaction.xa.*;
import com.ibm.db2.jdbc.app.*;

public class JTATxSuspend {

    public static void main(java.lang.String[] args) {
        JTATxSuspend test = new JTATxSuspend();

        test.setup();
        test.run();
    }

    /**
     * Handle the previous cleanup run so that this test can recommence.
     */
    public void setup() {

        Connection c = null;
        Statement s = null;
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            c = DriverManager.getConnection("jdbc:db2:*local");
            s = c.createStatement();

            try {
                s.executeUpdate("DROP TABLE CUJOSQL.JTATABLE");
            } catch (SQLException e) {
                // Ignore... doesn't exist
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.JTATABLE (COL1 CHAR (50))");
            s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('Fun with JTA')");
            s.executeUpdate("INSERT INTO CUJOSQL.JTATABLE VALUES('JTA is fun.')");

            s.close();
        } finally {
            if (c != null) {
                c.close();
            }
        }
    }

    /**
     * This test uses JTA support to handle transactions.
     */
    public void run() {
```

```

Connection c = null;

try {
    Context ctx = new InitialContext();

    // Assume the data source is backed by a UDBXADDataSource.
    UDBXADDataSource ds = (UDBXADDataSource) ctx.lookup("XADataSource");

    // From the DataSource, obtain an XAConnection object that
    // contains an XAResource and a Connection object.
    XAConnection xaConn = ds.getXAConnection();
    XAResource xaRes = xaConn.getXAResource();
    Connection c = xaConn.getConnection();

    // For XA transactions, a transaction identifier is required.
    // An implementation of the XID interface is not included with
    // the JDBC driver. Transactions with JTA for a
    // description of this interface to build a class for it.
    Xid xid = new XidImpl();

    // The connection from the XAResource can be used as any other
    // JDBC connection.
    Statement stmt = c.createStatement();

    // The XA resource must be notified before starting any
    // transactional work.
    xaRes.start(xid, XAResource.TMNOFLAGS);

    // Create a ResultSet during JDBC processing and fetch a row.
    ResultSet rs = stmt.executeUpdate("SELECT * FROM CUJOSQL.JTATABLE");
    rs.next();

    // The end method is called with the suspend option. The
    // ResultSets associated with the current transaction are 'on hold'.
    // They are neither gone nor accessible in this state.
    xaRes.end(xid, XAResource.TMSUSPEND);

    // Other work can be performed with the transaction.
    // As an example, you can create a statement and process a query.
    // This work and any other transactional work that the transaction may
    // perform is separate from the work done previously under the XID.
    Statement nonXASmt = conn.createStatement();
    ResultSet nonXARS = nonXASmt.executeQuery("SELECT * FROM CUJOSQL.JTATABLE");
    while (nonXARS.next()) {
        // Process here...
    }
    nonXARS.close();
    nonXASmt.close();

    // If an attempt is made to use any suspended transactions
    // resources, an exception results.
    try {
        rs.getString(1);
        System.out.println("Value of the first row is " + rs.getString(1));
    }
}

```

```

    } catch (SQLException e) {
        System.out.println("This was an expected exception - suspended ResultSet was used.");
    }

    // Resume the suspended transaction and complete the work on it.
    // The ResultSet is exactly as it was before the suspension.
    xaRes.start(newXid, XAResource.TMRESUME);
    rs.next();
    System.out.println("Value of the second row is " + rs.getString(1));

    // When the transaction has completed, end it
    // and commit any work under it.
    xaRes.end(xid, XAResource.TMNOFLAGS);
    int rc = xaRes.prepare(xid);
    xaRes.commit(xid, false);

} catch (Exception e) {
    System.out.println("Something has gone wrong.");
    e.printStackTrace();
} finally {
    try {
        if (c != null)
            c.close();
    } catch (SQLException e) {
        System.out.println("Note: Cleanup exception.");
        e.printStackTrace();
    }
}
}
}
}
}
}

```



Statement types

➤ The Statement interface and its PreparedStatement and CallableStatement subclasses are used to process structured query language (SQL) commands against the database. SQL statements cause the generation of ResultSet objects.

Subclasses of the Statement interface are created with a number of methods on the Connection interface. A single Connection object can have many Statement objects created under it simultaneously. In past releases, it was possible to give exact numbers of Statement objects that could be created. It is impossible to do so in this release because different types of Statement objects take different numbers of "handles" within the database engine. Therefore, the types of Statement objects you are using influence the number of statements that can be active under a connection at a single time.

An application calls the Statement.close method to indicate that the application has finished processing a statement. All Statement objects are closed when the connection that created them is closed. However, you should not fully rely on this behavior to close Statement objects. For example, if your application changes so that a connection pool is used instead of explicitly closing the connections, the application "leaks" statement handles because the connections never close. Closing Statement objects as soon as they are no longer required allows external database resources that the statement is using to be released immediately.

The native JDBC driver attempts to detect statement leaks and handles them on your behalf. However, relying on that support results in poorer performance.

To use statements and its subclasses, see the following:

Statements

A Statement object is used for processing a static SQL statement and obtaining the results produced by it.

PreparedStatement

PreparedStatement is a subclass of the Statement interface and provides support for adding parameters to SQL statements.

CallableStatements

CallableStatements extend the PreparedStatement interface and provide support for output and input/output parameters in addition to the input parameter support provided by PreparedStatement.

Due to the inheritance hierarchy that CallableStatement extends PreparedStatement which extends Statement, features of each interface are available in the class that extend the interface. For example, features of the Statement class are also supported in the PreparedStatement and CallableStatement classes. The main exception is the executeQuery, executeUpdate, and execute methods on the Statement class. These methods take in an SQL statement to dynamically process and cause exceptions if you attempt to use them with PreparedStatement or CallableStatement objects. <<

Statements: >> A Statement object is used for processing a static SQL statement and obtaining the results produced by it. Only one ResultSet for each Statement object can be open at a time. All statement methods that process an SQL statement implicitly close a statement's current ResultSet if an open one exists.

Create statements: Statement objects are created from Connection objects with the createStatement method. For example, assuming a Connection object named conn already exists, the following line of code creates a Statement object for passing SQL statements to the database:

```
Statement stmt = conn.createStatement();
```

<H7>Specify ResultSet characteristics

The characteristics of ResultSets are associated with the statement that eventually creates them. The Connection.createStatement method allows you to specify these ResultSet characteristics. The following are some examples of valid calls to the createStatement method:

Example: The createStatement method

Note: Read the Code example disclaimer for important legal information.

```
// The following is new in JDBC 2.0

Statement stmt2 = conn.createStatement(ResultSet.TYPE_SCROLL_INSENSITIVE,
ResultSet.CONCUR_UPDATEABLE);

// The following is new in JDBC 3.0

Statement stmt3 = conn.createStatement(ResultSet.TYPE_SCROLL_INSENSITIVE,
ResultSet.CONCUR_READ_ONLY, ResultSet.HOLD_CURSOR_OVER_COMMIT);
```

For more information about these characteristics, see ResultSets.

<H7>Process statements

Processing SQL statements with a Statement object is accomplished with the executeQuery(), executeUpdate(), and execute() methods.

Return results from SQL queries: If an SQL query statement returning a ResultSet object is to be processed, the executeQuery() method should be used. You can refer to the example program that uses a Statement object's executeQuery method to obtain a ResultSet.

Note: If an SQL statement processed with executeQuery does not return a ResultSet, an SQLException is thrown.

<h7>Return update counts for SQL Statements

If the SQL is known to be a Data Definition Language (DDL) statement or a Data Manipulation Language (DML) statement returning an update count, the executeUpdate() method should be used. The StatementExample program uses a Statement object's executeUpdate method.

<h7>Process SQL statements where the expected return is unknown

If the SQL statement type is not known, the execute method should be used. Once this method has been processed, the JDBC driver can tell the application what types of results the SQL statement has generated through API calls. The execute method returns true if the result is at least one ResultSet and false if the return value is an update count. Given this information, applications can use the statement method's getUpdateCount or getResultSet to retrieve the return value from processing the SQL statement. The StatementExecute program uses the execute method on a Statement object. This program expects a parameter to be passed that is an SQL statement. Without looking at the text of the SQL that you provide, the program processes the statement and determines information about what was processed.

Note: Calling the getUpdateCount method when the result is a ResultSet returns -1. Calling the getResultSet method when the result is an update count returns null.

The cancel method: The methods of the native JDBC driver are synchronized to prevent two threads running against the same object from corrupting the object. An exception is the cancel method. The cancel method can be used by one thread to stop a long running SQL statement on another thread for the same object. The native JDBC driver cannot force the thread to stop doing work; it can only request that it stop whatever task it was doing. For this reason, it still takes time for a cancelled statement to stop. The cancel method can be used to halt runaway SQL queries on the system. <<

Example: Use the Statement object's executeUpdate method: >> This is an example of how to use the Statement object's executeUpdate method.

Example: Use the Statement object's executeUpdate method

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import java.util.Properties;

public class StatementExample {

    public static void main(java.lang.String[] args)
    {

        // Suggestion: Load these from a properties object.
        String DRIVER = "com.ibm.db2.jdbc.app.DB2Driver";
        String URL     = "jdbc:db2://*local";

        // Register the native JDBC driver. If the driver cannot be
        // registered, the test cannot continue.
        try {
```

```

    Class.forName(DRIVER);
} catch (Exception e) {
    System.out.println("Driver failed to register.");
    System.out.println(e.getMessage());
    System.exit(1);
}

Connection c = null;
Statement s = null;

try {
    // Create the connection properties.
    Properties properties = new Properties ();
    properties.put ("user", "userid");
    properties.put ("password", "password");

    // Connect to the local iSeries database.
    c = DriverManager.getConnection(URL, properties);

    // Create a Statement object.
    s = c.createStatement();
    // Delete the test table if it exists. Note: This
    // example assumes that the collection MYLIBRARY
    // exists on the system.
    try {
        s.executeUpdate("DROP TABLE MYLIBRARY.MYTABLE");
    } catch (SQLException e) {
        // Just continue... the table probably does not exist.
    }

    // Run an SQL statement that creates a table in the database.
    s.executeUpdate("CREATE TABLE MYLIBRARY.MYTABLE (NAME VARCHAR(20), ID INTEGER)");

    // Run some SQL statements that insert records into the table.
    s.executeUpdate("INSERT INTO MYLIBRARY.MYTABLE (NAME, ID) VALUES ('RICH', 123)");
    s.executeUpdate("INSERT INTO MYLIBRARY.MYTABLE (NAME, ID) VALUES ('FRED', 456)");
    s.executeUpdate("INSERT INTO MYLIBRARY.MYTABLE (NAME, ID) VALUES ('MARK', 789)");

    // Run an SQL query on the table.
    ResultSet rs = s.executeQuery("SELECT * FROM MYLIBRARY.MYTABLE");

    // Display all the data in the table.
    while (rs.next()) {
        System.out.println("Employee " + rs.getString(1) + " has ID " + rs.getInt(2));
    }

} catch (SQLException sqle) {
    System.out.println("Database processing has failed.");
    System.out.println("Reason: " + sqle.getMessage());
} finally {
    // Close database resources
    try {
        try {
            if (s != null) {
                s.close();
            }
        } catch (SQLException e) {
            System.out.println("Cleanup failed to close Statement.");
        }
    }

    try {
        if (c != null) {
            c.close();
        }
    } catch (SQLException e) {
        System.out.println("Cleanup failed to close Connection.");
    }
}

```

```
}  
}  
}
```

PreparedStatements: ➤ PreparedStatements extend the Statement interface and provide support for adding parameters to SQL statements.

SQL statements that are passed to the database go through a two-step process in returning results to you. They are first prepared and then are processed. With Statement objects, these two phases appear to be one phase to your applications. PreparedStatements allow these two steps to be broken apart. The preparation step occurs when the object is created and the processing step occurs when the executeQuery, executeUpdate, or execute method are called on the PreparedStatement object.

Being able to split the SQL processing into separate phases are meaningless without the addition of parameter markers. Parameter markers are placed in an application so that it can tell the database that it does not have a specific value at preparation time, but that it provides one before processing time. Parameter markers are represented in SQL statements by question marks.

Parameter markers make it possible to make general SQL statements that are used for specific requests. For example, take the following SQL query statement:

```
SELECT * FROM EMPLOYEE_TABLE WHERE LASTNAME = 'DETTINGER'
```

This is a specific SQL statement that returns only one value; that is, information about an employee named Dettinger. By adding a parameter marker, the statement can become more flexible:

```
SELECT * FROM EMPLOYEE_TABLE WHERE LASTNAME = ?
```

By simply setting the parameter marker to a value, information can be obtained about any employee in the table.

PreparedStatements provide significant performance improvements over Statements because the previous Statement example can go through the preparation phase only once and then be processed repeatedly with different values for the parameter.

Note: Using PreparedStatements is a requirement to support the native JDBC driver's statement pooling.

Create PreparedStatement: The prepareStatement method is used to create new PreparedStatement objects. Unlike the createStatement method, the SQL statement must be supplied when the PreparedStatement object is created. At that time, the SQL statement is precompiled for use. For example, assuming a Connection object named conn already exists, the following example creates a PreparedStatement object and prepares the SQL statement for processing within the database.

```
PreparedStatement ps = conn.prepareStatement("SELECT * FROM EMPLOYEE_TABLE  
WHERE LASTNAME = ?");
```

Specify ResultSet characteristics and auto-generated key support: As with the createStatement method, the prepareStatement method is overloaded to provide support for specifying ResultSet characteristics. The prepareStatement method also has variations for working with auto-generated keys. The following are some examples of valid calls to the prepareStatement method:

Example: The prepareStatement method

Note: Read the Code example disclaimer for important legal information.

```
// New in JDBC 2.0  
  
PreparedStatement ps2 = conn.prepareStatement("SELECT * FROM  
EMPLOYEE_TABLE WHERE LASTNAME = ?",
```

```

ResultSet.TYPE_SCROLL_INSENSITIVE,
ResultSet.CONCUR_UPDATEABLE);

// New in JDBC 3.0

PreparedStatement ps3 = conn.prepareStatement("SELECT * FROM
EMPLOYEE_TABLE WHERE LASTNAME = ?",
ResultSet.TYPE_SCROLL_INSENSITIVE, ResultSet.CONCUR_UPDATEABLE,
ResultSet.HOLD_CURSOR_OVER_COMMIT);

PreparedStatement ps4 = conn.prepareStatement("SELECT * FROM
EMPLOYEE_TABLE WHERE LASTNAME = ?", Statement.RETURN_GENERATED_KEYS);

```

Handle parameters: Before a PreparedStatement object can be processed, each of the parameter markers must be set to some value. The PreparedStatement object provides a number of methods for setting parameters. All methods are of the form set<Type>, where <Type> is a Java data type. Some examples of these methods include setInt, setLong, setString, setTimestamp, setNull, and setBlob. Nearly all of these methods take two parameters:

- The first parameter is the index of the parameter within the statement. Parameter markers are numbered, starting with 1.
- The second parameter is the value to set the parameter to. There are a couple set<Type> methods that have additional parameters such as the length parameter on setBinaryStream.

Consult the Javadoc for the java.sql package for more information. Given the prepared SQL statement in the previous examples for the ps object, the following code illustrates how the parameter value is specified before processing:

```
ps.setString(1, 'Dettinger');
```

If an attempt is made to process a PreparedStatement with parameter markers that have not been set, an SQLException is thrown.

Note: Once set, parameter markers hold the same value between processes unless the following situations occur:

- The value is changed by another call to a set method.
- The value is removed when the clearParameters method is called.

The clearParameters method flags all parameters as being unset. After the call to clearParameters has been made, all the parameters must have the set method called again before the next process.

ParameterMetaData support: A new ParameterMetaData interface allows you to retrieve information about a parameter. This support is the compliment to ResultSetMetaData and is similar. Information such as the precision, scale, data type, data type name, and whether the parameter allows the null value are all provided.

See Example: ParameterMetaData on how to use this new support in an application program. <<

Process PreparedStatements: >> Processing SQL statements with a PreparedStatement object is accomplished with the executeQuery, executeUpdate, and execute methods like Statement objects are processed. Unlike Statement versions, no parameters are passed on these methods since the SQL statement was already provided when the object was created. Because PreparedStatement extends Statement, applications can attempt to call versions of executeQuery, executeUpdate, and execute methods that take a SQL statement. Doing so results in an SQLException being thrown.

Return results from SQL queries: If an SQL query statement that returns a ResultSet object is to be processed, the executeQuery method should be used. The PreparedStatementExample program uses a PreparedStatement object's executeQuery method to obtain a ResultSet.

Note: If an SQL statement processed with the `executeQuery` method does not return a `ResultSet`, an `SQLException` is thrown.

Return update counts for SQL statements: If the SQL is known to be a Data Definition Language (DDL) statement or a Data Manipulation Language (DML) statement that returns an update count, the `executeUpdate` method should be used. The `PreparedStatementExample` sample program uses a `PreparedStatement` object's `executeUpdate` method.

Process SQL statements where the expected return is unknown: If the SQL statement type is not known, the `execute` method should be used. Once this method has been processed, the JDBC driver can tell the application what results types the SQL statement generated through API calls. The `execute` method returns `true` if the result is at least one `ResultSet` and `false` if the return value is an update count. Given this information, applications can use the `getUpdateCount` or `getResultSet` statement methods to retrieve the return value from processing the SQL statement.

Note: Calling the `getUpdateCount` method when the result is a `ResultSet` returns `-1`. Calling the `getResultSet` method when the result is an update count returns `null`. <<

Example: Use `PreparedStatement` to obtain a `ResultSet`: >> This is an example of using a `PreparedStatement` object's `executeQuery` method to obtain a `ResultSet`.

Example: Use `PreparedStatement` to obtain a `ResultSet`

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import java.util.Properties;

public class PreparedStatementExample {

    public static void main(java.lang.String[] args)
    {
        // Load the following from a properties object.
        String DRIVER = "com.ibm.db2.jdbc.app.DB2Driver";
        String URL    = "jdbc:db2://*local";

        // Register the native JDBC driver. If the driver cannot
        // be registered, the test cannot continue.
        try {
            Class.forName(DRIVER);
        } catch (Exception e) {
            System.out.println("Driver failed to register.");
            System.out.println(e.getMessage());
            System.exit(1);
        }

        Connection c = null;
        Statement s = null;

        // This program creates a table that is
        // used by prepared statements later.
        try {
            // Create the connection properties.
            Properties properties = new Properties ();
            properties.put ("user", "userid");
            properties.put ("password", "password");

            // Connect to the local iSeries database.
            c = DriverManager.getConnection(URL, properties);

            // Create a Statement object.
            s = c.createStatement();
            // Delete the test table if it exists. Note that
```

```

// this example assumes throughout that the collection
// MYLIBRARY exists on the system.
try {
    s.executeUpdate("DROP TABLE MYLIBRARY.MYTABLE");
} catch (SQLException e) {
    // Just continue... the table probably did not exist.
}

// Run an SQL statement that creates a table in the database.
s.executeUpdate("CREATE TABLE MYLIBRARY.MYTABLE (NAME VARCHAR(20), ID INTEGER)");

} catch (SQLException sqle) {
    System.out.println("Database processing has failed.");
    System.out.println("Reason: " + sqle.getMessage());
} finally {
    // Close database resources
    try {
        if (s != null) {
            s.close();
        }
    } catch (SQLException e) {
        System.out.println("Cleanup failed to close Statement.");
    }
}

// This program then uses a prepared statement to insert many
// rows into the database.
PreparedStatement ps = null;
String[] nameArray = {"Rich", "Fred", "Mark", "Scott", "Jason",
    "John", "Jessica", "Blair", "Erica", "Barb"};
try {
    // Create a PreparedStatement object that is used to insert data into the
    // table.
    ps = c.prepareStatement("INSERT INTO MYLIBRARY.MYTABLE (NAME, ID) VALUES (?, ?)");

    for (int i = 0; i < nameArray.length; i++) {
        ps.setString(1, nameArray[i]);    // Set the Name from our array.
        ps.setInt(2, i+1);                // Set the ID.
        ps.executeUpdate();
    }

} catch (SQLException sqle) {
    System.out.println("Database processing has failed.");
    System.out.println("Reason: " + sqle.getMessage());
} finally {
    // Close database resources
    try {
        if (ps != null) {
            ps.close();
        }
    } catch (SQLException e) {
        System.out.println("Cleanup failed to close Statement.");
    }
}

// Use a prepared statement to query the database
// table that has been created and return data from it. In
// this example, the parameter used is arbitrarily set to
// 5, meaning return all rows where the ID field is less than
// or equal to 5.
try {
    ps = c.prepareStatement("SELECT * FROM MYLIBRARY.MYTABLE " +
        "WHERE ID <= ?");

    ps.setInt(1, 5);

```

```

        // Run an SQL query on the table.
        ResultSet rs = ps.executeQuery();
        // Display all the data in the table.
        while (rs.next()) {
            System.out.println("Employee " + rs.getString(1) + " has ID " + rs.getInt(2));
        }

    } catch (SQLException sqle) {
        System.out.println("Database processing has failed.");
        System.out.println("Reason: " + sqle.getMessage());
    } finally {
        // Close database resources
        try {
            if (ps != null) {
                ps.close();
            }
        } catch (SQLException e) {
            System.out.println("Cleanup failed to close Statement.");
        }

        try {
            if (c != null) {
                c.close();
            }
        } catch (SQLException e) {
            System.out.println("Cleanup failed to close Connection.");
        }
    }
}
}
}
}

```



Example: ParameterMetaData: ➤ This is an example of using the ParameterMetaData interface to retrieve information about parameters.

Example: ParameterMetaData

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////////////////////
//
// ParameterMetaData example. This program demonstrates
// the new support of JDBC 3.0 for learning information
// about parameters to a PreparedStatement.
//
// Command syntax:
//   java PMD
//
////////////////////////////////////////////////////
//
// This source is an example of the IBM Developer for Java JDBC driver.
// IBM grants you a nonexclusive license to use this as an example
// from which you can generate similar function tailored to
// your own specific needs.
//
// This sample code is provided by IBM for illustrative purposes
// only. These examples have not been thoroughly tested under all
// conditions. IBM, therefore, cannot guarantee or imply
// reliability, serviceability, or function of these programs.
//
// All programs contained herein are provided to you "AS IS"
// without any warranties of any kind. The implied warranties of
// merchantability and fitness for a particular purpose are

```



```

// expressly disclaimed.
//
// IBM Developer Kit for Java
// (C) Copyright IBM Corp. 2001
// All rights reserved.
// US Government Users Restricted Rights -
// Use, duplication, or disclosure restricted
// by GSA ADP Schedule Contract with IBM Corp.
//
//
//
import java.sql.*;

public class PMD {

    // Program entry point.
    public static void main(java.lang.String[] args)
    throws Exception
    {
        // Obtain setup.
        Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        PreparedStatement ps = c.prepareStatement("INSERT INTO CUJOSQL.MYTABLE VALUES(?, ?, ?)");
        ParameterMetaData pmd = ps.getParameterMetaData();

        for (int i = 1; i < pmd.getParameterCount(); i++) {
            System.out.println("Parameter number " + i);
            System.out.println("  Class name is " + pmd.getParameterClassName(i));
            // Note: Mode relates to input, output or inout
            System.out.println("  Mode is " + pmd.getParameterClassName(i));
            System.out.println("  Type is " + pmd.getParameterType(i));
            System.out.println("  Type name is " + pmd.getParameterTypeName(i));
            System.out.println("  Precision is " + pmd.getPrecision(i));
            System.out.println("  Scale is " + pmd.getScale(i));
            System.out.println("  Nullable? is " + pmd.isNullable(i));
            System.out.println("  Signed? is " + pmd.isSigned(i));
        }
    }
}

```



CallableStatements: ➤ The CallableStatement interface extends PreparedStatement and provides support for output and input/output parameters. The CallableStatement interface also has support for input parameters that is provided by the PreparedStatement interface.

The CallableStatement interface allows the use of SQL statements to call stored procedures. Stored procedures are programs that have a database interface. These programs possess the following:

- They can have input and output parameters, or parameters that are both input and output.
- They can have a return value.
- They have the ability to return multiple ResultSets.

Conceptually in JDBC, a stored procedure call is a single call to the database, but the program associated with the stored procedure may process hundreds of database requests. The stored procedure program may also perform a number of other programmatic tasks not typically done with SQL statements.

Because CallableStatements follow the PreparedStatement model of decoupling the preparation and processing phases, they have the potential for optimized reuse (see PreparedStatement for details). Since SQL statements of a stored procedure are bound into a program, they are processed as static SQL and further performance benefits can be gained that way. Encapsulating a lot of database work in a single,

reusable database call is an example of using stored procedures optimally. Only this call goes over the network to the other system, but the request can accomplish a lot of work on the remote system.

Create CallableStatements: The `prepareCall` method is used to create new `CallableStatement` objects. As with the `prepareStatement` method, the SQL statement must be supplied at the time that the `CallableStatement` object is created. At that time, the SQL statement is precompiled. For example, assuming a `Connection` object named `conn` already exists, the following creates a `CallableStatement` object and completes the preparation phase of getting the SQL statement ready for processing within the database:

```
PreparedStatement ps = conn.prepareStatement("? = CALL ADDEMPLOYEE(?, ?, ?)");
```

The `ADDEMPLOYEE` stored procedure takes input parameters for a new employee name, his social security number, and his manager's user ID. From this information, multiple company database tables may be updated with information about the employee such as his start date, division, department, and so on. Further, a stored procedure is a program that may generate standard user IDs and e-mail addresses for that employee. The stored procedure may also send an e-mail to the hiring manager with initial usernames and passwords; the hiring manager can then provide the information to the employee.

The `ADDEMPLOYEE` stored procedure is set up to have a return value. The return code may be a success or failure code that the calling program can use when a failure occurs. The return value may also be defined as the new employee's company ID number. Finally, the stored procedure program could have processed queries internally and have left the `ResultSets` from those queries open and available for the calling program. Querying all the new employee's information and making it available to the caller through a returned `ResultSet` is reasonable.

How to accomplish each of these types of tasks is covered in the following sections.

Specify ResultSet characteristics and auto-generated key support: As with `createStatement` and `prepareStatement`, there are multiple versions of `prepareCall` that provide support for specifying `ResultSet` characteristics. Unlike `prepareStatement`, the `prepareCall` method does not provide variations for working with auto-generated keys from `CallableStatements` (JDBC 3.0 does not support this concept.) The following are some examples of valid calls to the `prepareCall` method:

Example: The `prepareCall` method

```
// The following is new in JDBC 2.0
```

```
CallableStatement cs2 = conn.prepareCall("? = CALL ADDEMPLOYEE(?, ?, ?)",  
    ResultSet.TYPE_SCROLL_INSENSITIVE, ResultSet.CONCUR_UPDATEABLE);
```

```
// New in JDBC 3.0
```

```
CallableStatement cs3 = conn.prepareCall("? = CALL ADDEMPLOYEE(?, ?, ?)",  
    ResultSet.TYPE_SCROLL_INSENSITIVE, ResultSet.CONCUR_UPDATEABLE,  
    ResultSet.HOLD_CURSOR_OVER_COMMIT);
```

Handle parameters: As stated, `CallableStatement` objects may take three types of parameters:

- **IN**

IN parameters are handled in the same manner as `PreparedStatement`s. The various set methods of the inherited `PreparedStatement` class are used to set the parameters.

- **OUT**

OUT parameters are handled with the `registerOutParameter` method. The most common form of `registerOutParameter` takes an index parameter as the first parameter and an SQL type as the second parameter. This tells the JDBC driver what to expect for data from the parameter when the statement is processed. There are two other variations on the `registerOutParameter` method that can be found in the `java.sql` package Javadoc.

- **INOUT**

INOUT parameters require that the work for both IN parameters and OUT parameters be done. For

each INOUT parameter, you must call a set method and the registerOutParameter method before the statement can be processed. Failing to set or register any parameter results in an SQLException being thrown when the statement is processed.

Refer to Example: Create a procedure with input and output parameters for more information.

As with PreparedStatement, CallableStatement parameter values remain the same between processes unless you call a set method again. The clearParameters method does not affect parameters that are registered for output. After calling clearParameters, all IN parameters must be set to a value again, but all OUT parameters do not have to be registered again.

Note: The concept of parameters must not be confused with the index of a parameter marker. A stored procedure call expects a certain number of parameters that are passed to it. A particular SQL statement has ? characters (parameter markers) in it to represent values that are supplied at runtime. Consider the following example to see the difference between the two concepts:

```
CallableStatement cs = con.prepareCall("CALL PROC(?, \"SECOND\", ?)");  
  
cs.setString(1, "First");    //Parameter marker 1, Stored procedure parm 1  
  
cs.setString(2, "Third");    //Parameter marker 2, Stored procedure parm 3
```

Access stored procedure parameters by name: Parameters to stored procedures have names associated with them as the following stored procedure declaration shows:

Example: Stored procedure parameters

Note: Read the Code example disclaimer for important legal information.

```
CREATE  
PROCEDURE MYLIBRARY.APROC  
    (IN PARM1 INTEGER)  
LANGUAGE SQL SPECIFIC MYLIBRARY.APROC  
BODY: BEGIN  
    <Perform a task here...>  
END BODY
```

There is a single integer parameter with the name PARM1. In JDBC 3.0, there is support for specifying stored procedure parameters by name as well as by index. The code to set up a CallableStatement for this procedure is as follows:

```
CallableStatement cs = con.prepareCall("CALL APROC(?)");  
  
cs.setString("PARM1", 6);    //Sets input parameter at index 1 (PARM1) to 6.
```

For more information, see Process CallableStatements. <<

Process CallableStatements: >> Processing SQL stored procedure calls with a CallableStatement object is accomplished with the same methods that are used with a PreparedStatement object.

Return results for stored procedures: If an SQL query statement is processed within a stored procedure, the query results can be made available to the program calling the stored procedure. Multiple queries can also be called within the stored procedure and the calling program can process all the ResultSets that are available.

See Example: Create a procedure with multiple ResultSets for more information.

Note: If a stored procedure is processed with executeQuery and it does not return a ResultSet, an SQLException is thrown.

Access ResultSets concurrently: Return results for stored procedures deals with ResultSets and stored procedures and provides an example that works with all Java^(TM) Development Kit (JDK) releases. In the example, the ResultSets are processed in order from the first ResultSet that the stored procedure opened to the last ResultSet opened. One ResultSet is closed before the next is used.

In JDK 1.4, there is support for working with ResultSets from stored procedures concurrently.

Note: This feature was added to the underlying system support through the Command Line Interface (CLI) in V5R2. As a result, JDK 1.4 running on a system before V5R2 does not have this support available to it.

Return update counts for stored procedures: Returning update counts for stored procedures is a feature discussed in the JDBC specification, but it is not currently supported on the iSeries platform. There is no way to return multiple update counts from a stored procedure call. If an update count is needed from a processed SQL statement within a stored procedure, there are two ways of returning the value:

- Return the value as an output parameter.
- Pass back the value as the return value from the parameter. This is a special case of an output parameter. See Process stored procedures that have a return value for more information.

Process stored procedures where the expected return is unknown: If the results from a stored procedure call are not known, the execute method should be used. Once this method has been processed, the JDBC driver can tell the application what types of results the stored procedure generated through API calls. The execute method returns true if the result is one or more ResultSets. Updating counts do not come from stored procedure calls.

Process stored procedures that have a return value: The iSeries platform supports stored procedures that have a return value similar to a function's return value. The return value from a stored procedure is labeled like other parameters marks and is labeled such that it is assigned by the stored procedure call. An example of this is as follows:

```
? = CALL MYPROC(?, ?, ?)
```

The return value from a stored procedure call is always an integer type and must be registered like any other output parameter.

See Example: Create a procedure with return values for more information. [⏪](#)

Example: CallableStatement interface for IBM Developer Kit for Java: This is an example of how to use the CallableStatement interface.

Example: CallableStatement interface

Note: Read the Code example disclaimer for important legal information.

```
// Connect to iSeries server.
Connection c = DriverManager.getConnection("jdbc:db2://mySystem");

// Create the CallableStatement object.
// It precompiles the specified call to a stored procedure.
// The question marks indicate where input parameters must be set and
// where output parameters can be retrieved.
// The first two parameters are input parameters, and the third parameter is an output parameter.
CallableStatement cs = c.prepareCall("CALL MYLIBRARY.ADD (?, ?, ?)");

// Set input parameters.
cs.setInt (1, 123);
cs.setInt (2, 234);

// Register the type of the output parameter.
cs.registerOutParameter (3, Types.INTEGER);
```

```

// Run the stored procedure.
cs.execute ();

// Get the value of the output parameter.
int sum = cs.getInt (3);

// Close the CallableStatement and the Connection.
cs.close();
c.close();

```

For more information, see [CallableStatements](#).

Example: Create a procedure with multiple ResultSets: ➤ **Note:** Read the Code example disclaimer for important legal information. ⏪

```

import java.sql.*;
import java.util.Properties;

public class CallableStatementExample1 {

    public static void main(java.lang.String[] args) {

        // Register the Native JDBC driver. If we cannot
        // register the driver, the test cannot continue.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

            // Create the connection properties
            Properties properties = new Properties ();
            properties.put ("user", "userid");
            properties.put ("password", "password");

            // Connect to the local iSeries database
            Connection c = DriverManager.getConnection("jdbc:db2://*local", properties);

            Statement s = c.createStatement();

            // Create a procedure with multiple ResultSets.
            String sql = "CREATE PROCEDURE MYLIBRARY.SQLSPEX1 " +
                "RESULT SET 2 LANGUAGE SQL READS SQL DATA SPECIFIC MYLIBRARY.SQLSPEX1 " +
                "EX1: BEGIN " +
                "    DECLARE C1 CURSOR FOR SELECT * FROM QSYS2.SYSPROCS " +
                "        WHERE SPECIFIC_SCHEMA = 'MYLIBRARY'; " +
                "    DECLARE C2 CURSOR FOR SELECT * FROM QSYS2.SYSPARMS " +
                "        WHERE SPECIFIC_SCHEMA = 'MYLIBRARY'; " +
                "    OPEN C1; " +
                "    OPEN C2; " +
                "    SET RESULT SETS CURSOR C1, CURSOR C2; " +
                "END EX1 ";

            try {
                s.executeUpdate(sql);
            } catch (SQLException e) {
                // NOTE: We are ignoring the error here. We are making
                // the assumption that the only reason this fails
                // is because the procedure already exists. Other
                // reasons that it could fail are because the C compiler
                // is not found to compile the procedure or because
                // collection MYLIBRARY does not exist on the system.
            }
            s.close();

            // Now use JDBC to run the procedure and get the results back. In
            // this case we are going to get information about 'MYLIBRARY's stored
            // procedures (which is also where we created this procedure, thereby
            // ensuring that there is something to get.

```

```

CallableStatement cs = c.prepareCall("CALL MYLIBRARY.SQLSPEX1");

ResultSet rs = cs.executeQuery();

// We now have the first ResultSet object that the stored procedure
// left open. Use it.
int i = 1;
while (rs.next()) {
    System.out.println("MYLIBRARY stored procedure " + i + " is " + rs.getString(1) + "." +
        rs.getString(2));
    i++;
}
System.out.println("");

// Now get the next ResultSet object from the system - the previous
// one is automatically closed.
if (!cs.getMoreResults()) {
    System.out.println("Something went wrong. There should have been another ResultSet, exiting.");
    System.exit(0);
}
rs = cs.getResultSet();

// We now have the second ResultSet object that the stored procedure
// left open. Use that one.
i = 1;
while (rs.next()) {
    System.out.println("MYLIBRARY procedure " + rs.getString(1) + "." + rs.getString(2) +
        " parameter: " + rs.getInt(3) + " direction: " + rs.getString(4) +
        " data type: " + rs.getString(5));
    i++;
}

if (i == 1) {
    System.out.println("None of the stored procedures have any parameters.");
}

if (cs.getMoreResults()) {
    System.out.println("Something went wrong, there should not be another ResultSet.");
    System.exit(0);
}

cs.close(); // close the CallableStatement object
c.close(); // close the Connection object.

} catch (Exception e) {
    System.out.println("Something failed..");
    System.out.println("Reason: " + e.getMessage());
    e.printStackTrace();
}
}
}

```

Example: Create a procedure with input and output parameters: **Note:** Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import java.util.Properties;

public class CallableStatementExample2 {

    public static void main(java.lang.String[] args) {

        // Register the Native JDBC driver. If we cannot

```

```

// register the driver, the test cannot continue.
try {
    Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

    // Create the connection properties
    Properties properties = new Properties ();
    properties.put ("user", "userid");
    properties.put ("password", "password");

    // Connect to the local iSeries database
    Connection c = DriverManager.getConnection("jdbc:db2://*local", properties);

    Statement s = c.createStatement();

    // Create a procedure with in, out, and in/out parameters.
    String sql = "CREATE PROCEDURE MYLIBRARY.SQLSPEX2 " +
        "(IN P1 INTEGER, OUT P2 INTEGER, INOUT P3 INTEGER) " +
        "LANGUAGE SQL SPECIFIC MYLIBRARY.SQLSPEX2 " +
        "EX2: BEGIN " +
        "    SET P2 = P1 + 1; " +
        "    SET P3 = P3 + 1; " +
        "END EX2 ";

    try {
        s.executeUpdate(sql);
    } catch (SQLException e) {
        // NOTE: We are ignoring the error here. We are making
        // the assumption that the only reason this fails
        // is because the procedure already exists. Other
        // reasons that it could fail are because the C compiler
        // is not found to compile the procedure or because
        // collection MYLIBRARY does not exist on the system.
    }
    s.close();

    // Prepare a callable statement used to run the procedure.
    CallableStatement cs = c.prepareCall("CALL MYLIBRARY.SQLSPEX2(?, ?, ?)");

    // All input parameters must be set and all output parameters must
    // be registered. Notice that this means we have two calls to make
    // for an input output parameter.
    cs.setInt(1, 5);
    cs.setInt(3, 10);
    cs.registerOutParameter(2, Types.INTEGER);
    cs.registerOutParameter(3, Types.INTEGER);

    // Run the procedure
    cs.executeUpdate();

    // Verify the output parameters have the desired values.
    System.out.println("The value of P2 should be P1 (5) + 1 = 6. --> " + cs.getInt(2));
    System.out.println("The value of P3 should be P3 (10) + 1 = 11. --> " + cs.getInt(3));

    cs.close(); // close the CallableStatement object
    c.close(); // close the Connection object.

} catch (Exception e) {
    System.out.println("Something failed..");
    System.out.println("Reason: " + e.getMessage());
    e.printStackTrace();
}
}
}

```

Example: Create a procedure with return values: Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import java.util.Properties;

public class CallableStatementExample3 {

    public static void main(java.lang.String[] args) {

        // Register the native JDBC driver. If the driver cannot
        // be registered, the test cannot continue.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

            // Create the connection properties
            Properties properties = new Properties ();
            properties.put ("user", "userid");
            properties.put ("password", "password");

            // Connect to the local iSeries database
            Connection c = DriverManager.getConnection("jdbc:db2://*local", properties);

            Statement s = c.createStatement();

            // Create a procedure with a return value. Note that return value support
            // is new in V4R5.
            String sql = "CREATE PROCEDURE MYLIBRARY.SQLSPEX3 " +
                " LANGUAGE SQL SPECIFIC MYLIBRARY.SQLSPEX3 " +
                " EX3: BEGIN " +
                "     RETURN 1976; " +
                " END EX3 ";

            try {
                s.executeUpdate(sql);
            } catch (SQLException e) {
                // NOTE: The error is ignored here. The assumptions is
                //     made that the only reason this fails is
                //     because the procedure already exists. Other
                //     reasons that it could fail are because the C compiler
                //     is not found to compile the procedure or because
                //     collection MYLIBRARY does not exist on the system.
            }
            s.close();

            // Prepare a callable statement used to run the procedure.
            CallableStatement cs = c.prepareCall("? = CALL MYLIBRARY.SQLSPEX3");

            // You still need to register the output parameter.
            cs.registerOutParameter(1, Types.INTEGER);

            // Run the procedure.
            cs.executeUpdate();

            // Show that the correct value is returned.
            System.out.println("The return value should always be 1976 for this example: --> " + cs.getInt(1));

            cs.close(); // close the CallableStatement object
            c.close(); // close the Connection object.

        } catch (Exception e) {
            System.out.println("Something failed..");
            System.out.println("Reason: " + e.getMessage());
            e.printStackTrace();
        }
    }
}

```


ResultSets

» The ResultSet interface provides access to the results generated by running queries. Conceptually, data of a ResultSet can be thought of as a table with a specific number of columns and a specific number of rows. By default, the table rows are retrieved in sequence. Within a row, column values can be accessed in any order.

To use the ResultSet object, see the following:

ResultSet characteristics

This section discusses ResultSet characteristics such as the following:

- ResultSet types
- Concurrency
- Ability to close the ResultSet by committing the Connection object.
- Specification of ResultSet characteristics

Cursor movement

The iSeries Java^(TM) Database Connectivity (JDBC) drivers support scrollable ResultSets. With a scrollable ResultSet, you can process rows of data in any order using a number of cursor-positioning methods.

Retrieve ResultSet data

Find out how a ResultSet object provides methods for obtaining column data for a row.

Change ResultSets

With the iSeries JDBC drivers, you can change ResultSets by performing these tasks:

- Update rows
- Delete rows
- Insert rows
- Change positioned updates

Create ResultSets

You can create a ResultSet object by using the executeQuery methods provided by the Statement, PreparedStatement, or CallableStatement interfaces. This section also discusses closing ResultSet objects when they are no longer needed in your application.



ResultSet characteristics: » By default, all created ResultSets have a type of forward only, a concurrency of read only, and cursors are held over commit boundaries. An exception to this is that WebSphere currently changes the cursor holdability default so that cursors are implicitly closed when committed. These characteristics are configurable through methods that are accessible on Statement, PreparedStatement, and CallableStatement objects.

ResultSet types: The ResultSet type specifies the following about the ResultSet:

- Whether the ResultSet is scrollable.
- The types of Java^(TM) Database Connectivity (JDBC) ResultSets that are defined by constants on the ResultSet interface.

Definitions of these ResultSet types are as follows:

TYPE_FORWARD_ONLY

A cursor that can only be used to process from the beginning of a ResultSet to the end of it. This is the default type.

TYPE_SCROLL_INSENSITIVE

A cursor that can be used to scroll in various ways through a ResultSet. This type of cursor is insensitive to changes made to the database while it is open. It contains rows that satisfy the query when the query was processed or when data is fetched.

TYPE_SCROLL_SENSITIVE

A cursor that can be used to scroll in various ways through a ResultSet. This type of cursor is sensitive to changes made to the database while it is open. Changes to the database have a direct impact on the ResultSet data.

JDBC 1.0 ResultSets are always forward only. Scrollable cursors were added in JDBC 2.0.

Note: The blocking enabled and block size connection properties affect the degree of sensitivity of a TYPE_SCROLL_SENSITIVE cursor. Blocking enhances performance by caching data in the JDBC driver layer itself.

See Example: Sensitive and insensitive ResultSets that shows the difference between sensitive and insensitive ResultSets when rows are inserted into a table.

See Example: ResultSet sensitivity that shows how a change can affect a where clause of an SQL statement based on the sensitivity of the ResultSet.

Concurrency: Concurrency determines whether the ResultSet can be updated. The types are again defined by constants in the ResultSet interface. The available concurrency settings are as follows:

CONCUR_READ_ONLY

A ResultSet that can only be used for reading data out of the database. This is the default setting.

CONCUR_UPDATEABLE

A ResultSet that allows you to make changes to it. These changes can be placed into the underlying database. See Change ResultSets for more information.

JDBC 1.0 ResultSets are always forward only. Updateable ResultSets were added in JDBC 2.0.

Note: According to the JDBC specification, the JDBC driver is allowed to change the ResultSet type of the ResultSet concurrency setting if the values cannot be used together. In such cases, the JDBC driver places a warning on the Connection object.

There is one situation where the application specifies a TYPE_SCROLL_INSENSITIVE, CONCUR_UPDATEABLE ResultSet. Insensitivity is implemented in the database engine by making a copy of the data. You are then not allowed to make updates through that copy to the underlying database. If you specify this combination, the driver changes the sensitivity to TYPE_SCROLL_SENSITIVE and create the warning indicating that your request has been changed.

Holdability: The holdability characteristic determines whether calling commit on the Connection object closes the ResultSet. The JDBC API for working with the holdability characteristic is new in version 3.0. However, the native JDBC driver has provided a connection property for several releases that allows you to specify that default for all ResultSets created under the connection (see Connection properties and DataSource properties). The API support overrides any setting for the connection property. Values for the holdability characteristic are defined by ResultSet constants and are as follows:

HOLD_CURSOR_OVER_COMMIT

All open cursors remain open when the commit clause is called. This is the native JDBC default value.

CLOSE_CURSORS_ON_COMMIT

All open cursors are closed when commit clause is called.

Note: Calling rollback on a connection always closes all open cursors. This is a little known fact, but a common way for databases to handle cursors.

According to the JDBC specification, the default for cursor holdability is implementation-defined. Some platforms choose to use `CLOSE_CURSORS_ON_COMMIT` as the default. This does not usually become an issue for most applications, but you must be aware of what the driver you are working with does if you are working with cursors across commit boundaries. The Toolbox JDBC driver also uses the `HOLD_CURSORS_ON_COMMIT` default, but the JDBC driver for UDB for Windows^(R) NT has a default of `CLOSE_CURSORS_ON_COMMIT`.

Specify ResultSet characteristics: A `ResultSet`'s characteristics do not change once the `ResultSet` object has been created. Therefore, the characteristics have to be specified before creating the object. You can specify these characteristics through overloaded variations of the `createStatement`, `prepareStatement`, and `prepareCall` methods.

See the following topics to specify `ResultSet` characteristics:

- Specify `ResultSet` characteristics for `Statements`
- Specify `ResultSet` characteristics and automatically generated key support for `PreparedStatement`
- Specify `ResultSet` characteristics and auto-generated key support for `CallableStatements`

Note: There are `ResultSet` methods to obtain the `ResultSet` type and the concurrency of the `ResultSet`, but there is no method to obtain the holdability of the `ResultSet`. <<

Example: Sensitive and insensitive ResultSets: >> The following example shows the difference between sensitive and insensitive `ResultSets` when rows are inserted into a table.

Example: Sensitive and insensitive `ResultSets`

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;

public class Sensitive {

    public Connection connection = null;

    public static void main(java.lang.String[] args) {
        Sensitive test = new Sensitive();

        test.setup();
        test.run("sensitive");
        test.cleanup();

        test.setup();
        test.run("insensitive");
        test.cleanup();
    }

    public void setup() {

        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            connection = DriverManager.getConnection("jdbc:db2:*local");

            Statement s = connection.createStatement();
            try {
                s.executeUpdate("drop table cujosql.sensitive");
            } catch (SQLException e) {
                // Ignored.
            }
        }
    }
}
```

```

        s.executeUpdate("create table cujosql.sensitive(col1 int)");
        s.executeUpdate("insert into cujosql.sensitive values(1)");
        s.executeUpdate("insert into cujosql.sensitive values(2)");
        s.executeUpdate("insert into cujosql.sensitive values(3)");
        s.executeUpdate("insert into cujosql.sensitive values(4)");
        s.executeUpdate("insert into cujosql.sensitive values(5)");
        s.close();
    } catch (Exception e) {
        System.out.println("Caught exception: " + e.getMessage());
        if (e instanceof SQLException) {
            SQLException another = ((SQLException) e).getNextException();
            System.out.println("Another: " + another.getMessage());
        }
    }
}

public void run(String sensitivity) {
    try {
        Statement s = null;
        if (sensitivity.equalsIgnoreCase("insensitive")) {
            System.out.println("creating a TYPE_SCROLL_INSENSITIVE cursor");
            s = connection.createStatement(ResultSet.TYPE_SCROLL_INSENSITIVE, ResultSet.CONCUR_READ_ONLY);
        } else {
            System.out.println("creating a TYPE_SCROLL_SENSITIVE cursor");
            s = connection.createStatement(ResultSet.TYPE_SCROLL_SENSITIVE, ResultSet.CONCUR_READ_ONLY);
        }

        ResultSet rs = s.executeQuery("select * From cujosql.sensitive");

        // Fetch the five values that are there.
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        System.out.println("fetched the five rows...");

        // Note: If you fetch the last row, the ResultSet looks
        //        closed and subsequent new rows that are added
        //        are not be recognized.

        // Allow another statement to insert a new value.
        Statement s2 = connection.createStatement();
        s2.executeUpdate("insert into cujosql.sensitive values(6)");
        s2.close();

        // Whether a row is recognized is based on the sensitivity setting.
        if (rs.next()) {
            System.out.println("There is a row now: " + rs.getInt(1));
        } else {
            System.out.println("No more rows.");
        }
    } catch (SQLException e) {
        System.out.println("SQLException exception: ");
        System.out.println("Message:....." + e.getMessage());
        System.out.println("SQLState:...." + e.getSQLState());
        System.out.println("Vendor Code:.." + e.getErrorCode());
    }
}

```

```

        System.out.println("-----");
        e.printStackTrace();
    }
    catch (Exception ex) {
        System.out.println("An exception other than an SQLException was thrown: ");
        ex.printStackTrace();
    }
}

public void cleanup() {
    try {
        connection.close();
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}
}

```



Example: ResultSet sensitivity: ➤ The following example shows how a change can affect a where clause of an SQL statement based on the sensitivity of the ResultSet.

Example: ResultSet sensitivity

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;

public class Sensitive2 {

    public Connection connection = null;

    public static void main(java.lang.String[] args) {
        Sensitive2 test = new Sensitive2();

        test.setup();
        test.run("sensitive");
        test.cleanup();

        test.setup();
        test.run("insensitive");
        test.cleanup();
    }

    public void setup() {

        try {
            System.out.println("Native JDBC used");
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            connection = DriverManager.getConnection("jdbc:db2:*local");

            Statement s = connection.createStatement();
            try {
                s.executeUpdate("drop table cujosql.sensitive");
            } catch (SQLException e) {
                // Ignored.
            }

            s.executeUpdate("create table cujosql.sensitive(coll int)");
            s.executeUpdate("insert into cujosql.sensitive values(1)");
            s.executeUpdate("insert into cujosql.sensitive values(2)");
        }
    }
}

```

```

s.executeUpdate("insert into cujosql.sensitive values(3)");
s.executeUpdate("insert into cujosql.sensitive values(4)");
s.executeUpdate("insert into cujosql.sensitive values(5)");

try {
    s.executeUpdate("drop table cujosql.sensitive2");
} catch (SQLException e) {
    // Ignored.
}

s.executeUpdate("create table cujosql.sensitive2(col2 int)");
s.executeUpdate("insert into cujosql.sensitive2 values(1)");
s.executeUpdate("insert into cujosql.sensitive2 values(2)");
s.executeUpdate("insert into cujosql.sensitive2 values(3)");
s.executeUpdate("insert into cujosql.sensitive2 values(4)");
s.executeUpdate("insert into cujosql.sensitive2 values(5)");

s.close();

} catch (Exception e) {
    System.out.println("Caught exception: " + e.getMessage());
    if (e instanceof SQLException) {
        SQLException another = ((SQLException) e).getNextException();
        System.out.println("Another: " + another.getMessage());
    }
}
}

public void run(String sensitivity) {
    try {

        Statement s = null;
        if (sensitivity.equalsIgnoreCase("insensitive")) {
            System.out.println("creating a TYPE_SCROLL_INSENSITIVE cursor");
            s = connection.createStatement(ResultSet.TYPE_SCROLL_INSENSITIVE, ResultSet.CONCUR_READ_ONLY);
        } else {
            System.out.println("creating a TYPE_SCROLL_SENSITIVE cursor");
            s = connection.createStatement(ResultSet.TYPE_SCROLL_SENSITIVE, ResultSet.CONCUR_READ_ONLY);
        }

        ResultSet rs = s.executeQuery("select col1, col2 From cujosql.sensitive,
                                     cujosql.sensitive2 where col1 = col2");

        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));
        rs.next();
        System.out.println("value is " + rs.getInt(1));

        System.out.println("fetched the four rows...");

        // Another statement creates a value that does not fit the where clause.
        Statement s2 = connection.createStatement(ResultSet.TYPE_SCROLL_SENSITIVE, ResultSet.CONCUR_UPDATEABLE);
        ResultSet rs2 = s2.executeQuery("select * from cujosql.sensitive where col1 = 5 FOR UPDATE");
        rs2.next();
        rs2.updateInt(1, -1);
        rs2.updateRow();
        s2.close();

        if (rs.next()) {
            System.out.println("There is still a row: " + rs.getInt(1));
        } else {

```

```

        System.out.println("No more rows.");
    }

    } catch (SQLException e) {
        System.out.println("SQLException exception: ");
        System.out.println("Message:....." + e.getMessage());
        System.out.println("SQLState:...." + e.getSQLState());
        System.out.println("Vendor Code:." + e.getErrorCode());
        System.out.println("-----");
        e.printStackTrace();
    }
    catch (Exception ex) {
        System.out.println("An exception other than an SQLException was thrown: ");
        ex.printStackTrace();
    }
}

public void cleanup() {
    try {
        connection.close();
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}
}

```



Cursor movement: ➤ The `ResultSet.next` method is used to move through a `ResultSet` one row at a time. With Java^(TM) Database Connectivity (JDBC) 2.0, the iSeries JDBC drivers support scrollable `ResultSet`s. Scrollable `ResultSet`s allow processing the rows of data in any order by using the previous, absolute, relative, first, and last methods.

By default, JDBC `ResultSet`s are always forward only, meaning that the only valid cursor-positioning method to call is `next()`. You have to explicitly request a scrollable `ResultSet`. See `ResultSet` types for more information.

With a scrollable `ResultSet`, you can use the following cursor-positioning methods:

Method	Description
Next	This method moves the cursor forward one row in the <code>ResultSet</code> . The method returns true if the cursor is positioned on a valid row and false otherwise.
Previous	The method moves the cursor backward one row in the <code>ResultSet</code> . The method returns true if the cursor is positioned on a valid row and false otherwise.
First	The method moves the cursor to the first row in the <code>ResultSet</code> . The method returns true if the cursor is positioned on the first row and false if the <code>ResultSet</code> is empty.
Last	The method moves the cursor to the last row in the <code>ResultSet</code> . The method returns true if the cursor is positioned on the last row and false if the <code>ResultSet</code> is empty.

Method	Description
BeforeFirst	The method moves the cursor immediately before the first row in the ResultSet. For an empty ResultSet, this method has no effect. There is no return value from this method.
AfterLast	The method moves the cursor immediately after the last row in the ResultSet. For an empty ResultSet, this method has no effect. There is no return value from this method.
Relative (int rows)	The method moves the cursor relative to its current position. <ul style="list-style-type: none"> • If rows is 0, this method has no effect. • If rows is positive, the cursor is moved forward that many rows. If there are fewer rows between the current position and the end of the ResultSet than specified by the input parameters, this method operates like afterLast. • If rows is negative, the cursor is moved backward that many rows. If there are fewer rows between the current position and the end of the ResultSet than specified by the input parameter, this method operates like beforeFirst. <p>The method returns true if the cursor is positioned on a valid row and false otherwise.</p>
Absolute (int row)	The method moves the cursor to the row specified by row value. If row value is positive, the cursor is positioned that many rows from the beginning of the ResultSet. The first row is numbered 1, the second is 2, and so on. If there are fewer rows in the ResultSet than specified by the row value, this method operates the same way as afterLast. If row value is negative, the cursor is positioned that many rows from the end of the ResultSet. The last row is numbered -1, the second to last is -2, and so on. If there are fewer rows in the ResultSet than specified by the row value, this method operates the same way beforeFirst. If row value is 0, this method operates the same way as beforeFirst. The method returns true if the cursor is positioned on a valid row and false otherwise.



Retrieve ResultSet data: ➤ The ResultSet object provides several methods for obtaining column data for a row. All are of the form `get<Type>`, where `<Type>` is a Java^(TM) data type. Some examples of these methods include `getInt`, `getLong`, `getString`, `getTimestamp`, and `getBlob`. Nearly all of these methods take a single parameter that is either the column index within the ResultSet or the column name.

ResultSet columns are numbered, starting with 1. If the column name is used and there is more than one column in the ResultSet with the same name, the first one is returned. There are some `get<Type>` methods that have additional parameters, such as the optional Calendar object, which can be passed to `getTime`, `getDate`, and `getTimestamp`. Refer to the Javadoc for the `java.sql` package for full details.

For `get` methods that return objects, the return value is null when the column in the ResultSet is null. For primitive types, null cannot be returned. In these cases, the value is 0 or false. If an application must distinguish between null, and 0 or false, the `wasNull` method can be used immediately after the call. This method can then determine whether the value was an actual 0 or false value, or if that value was returned because the ResultSet value was indeed null.

See Example: ResultSet interface for IBM Developer Kit for Java for an example on how to use the ResultSet interface.

ResultSetMetaData support: When the `getMetaData` method is called on a `ResultSet` object, the method returns a `ResultSetMetaData` object describing the columns of that `ResultSet` object. When the SQL statement being processed is unknown until runtime, the `ResultSetMetaData` can be used to determine what `get` methods should be used to retrieve the data. The following code example uses `ResultSetMetaData` to determine each column type in the result set:

Example: Use `ResultSetMetaData` to determine each column type in a result set

Note: Read the Code example disclaimer for important legal information.

```
ResultSet rs = stmt.executeQuery(sqlString);
ResultSetMetaData rsmd = rs.getMetaData();
int colType [] = new int[rsmd.getColumnCount()];
for (int idx = 0, int col = 1; idx < colType.length; idx++, col++)
    colType[idx] = rsmd.getColumnType(col);
```

See Example: `ResultSetMetaData` interface for IBM Developer Kit for Java for an example of how to use the `ResultSetMetaData` interface. <<

Change ResultSets: >> The default setting for `ResultSets` is read only. However, with Java^(TM) Database Connectivity (JDBC) 2.0, the iSeries JDBC drivers provide complete support for updateable `ResultSets`. You can refer to `ResultSet` concurrency on how to update `ResultSets`.

Update rows: Rows may be updated in a database table through the `ResultSet` interface. The steps involved in this process are the following:

1. Change the values for a specific row using various `update<Type>` methods, where `<Type>` is a Java data type. These `update<Type>` methods correspond to the `get<Type>` methods available for retrieving values.
2. Apply the rows to the underlying database.

The database itself is not updated until the second step. Updating columns in a `ResultSet` without calling the `updateRow` method does not make any changes to the database.

Planned updates to a row can be thrown away with the `cancelUpdates` method. Once the `updateRow` method is called, changes to the database are final and cannot be undone.

Note: The `rowUpdated` method always returns false as the database does not have a way to point out which rows have been updated. Correspondingly, the `updatesAreDetected` method returns false.

Delete rows: Rows may be deleted in a database table through the `ResultSet` interface. The `deleteRow` method is provided and deletes the current row.

Insert rows: Rows may be inserted into a database table through the `ResultSet` interface. This process makes use of an "insert row" which applications specifically move the cursor to and build the values they want to insert into the database. The steps involved in this process are as follows:

1. Position the cursor on the insert row.
2. Set each of the values for the columns in the new row.
3. Insert the row into the database and optionally move the cursor back to the current row within the `ResultSet`.

Note: New rows are not inserted into the table where the cursor is positioned. They are typically added to the end of the table data space. A relational database is not position-dependent by default. For example, you should not expect to move the cursor to the third row and insert something that shows up before the fourth row when subsequent users fetch the data.

Support for positioned updates: Besides the method for updating the database through a ResultSet, SQL statements can be used to issue positioned updates. This support relies on using named cursors. JDBC provides the setCursorName method from Statement and the getCursorName method from ResultSet to provide access to these values.

Two DatabaseMetaData methods, supportsPositionedUpdate and supportsPositionedDelete, both return true as this feature is supported with the native JDBC driver.

See Example: Change values with a statement through another statement's cursor for more information.

See Example: Remove values from a table through another statement's cursor for more information. <<

Example: Remove values from a table through another statement's cursor: >> This is an example of how to remove values from a table through another statement's cursor.

Example: Remove values from a table through another statement's cursor

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;

public class UsingPositionedDelete {
    public Connection connection = null;
    public static void main(java.lang.String[] args) {
        UsingPositionedDelete test = new UsingPositionedDelete();

        test.setup();
        test.displayTable();

        test.run();
        test.displayTable();

        test.cleanup();
    }

    /**
    Handle all the required setup work.
    **/
    public void setup() {
        try {
            // Register the JDBC driver.
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

            connection = DriverManager.getConnection("jdbc:db2:*local");

            Statement s = connection.createStatement();
            try {
                s.executeUpdate("DROP TABLE CUJOSQL.WHERECUREX");
            } catch (SQLException e) {
                // Ignore problems here.
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.WHERECUREX ( " +
                "COL_IND INT, COL_VALUE CHAR(20)) ");

            for (int i = 1; i <= 10; i++) {
                s.executeUpdate("INSERT INTO CUJOSQL.WHERECUREX VALUES(" + i + ", 'FIRST')");
            }

            s.close();

        } catch (Exception e) {
```

```

        System.out.println("Caught exception: " + e.getMessage());
        e.printStackTrace();
    }
}

/**
In this section, all the code to perform the testing should
be added. If only one connection to the database is needed,
the global variable 'connection' can be used.
**/
public void run() {
    try {
        Statement stmt1 = connection.createStatement();

        // Update each value using next().
        stmt1.setCursorName("CUJO");
        ResultSet rs = stmt1.executeQuery ("SELECT * FROM CUJOSQL.WHERECUREX " +
            "FOR UPDATE OF COL_VALUE");

        System.out.println("Cursor name is " + rs.getCursorName());

        PreparedStatement stmt2 = connection.prepareStatement ("DELETE FROM "
            + " CUJOSQL.WHERECUREX WHERE CURRENT OF "
            + rs.getCursorName ());

        // Loop through the ResultSet and update every other entry.
        while (rs.next ()) {
            if (rs.next())
                stmt2.execute ();
        }

        // Clean up the resources after they have been used.
        rs.close ();
        stmt2.close ();

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

/**
In this section, put all clean-up work for testing.
**/
public void cleanup() {
    try {
        // Close the global connection opened in setup().
        connection.close();

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

/**
Display the contents of the table.
**/
public void displayTable()

```

```

{
    try {
        Statement s = connection.createStatement();
        ResultSet rs = s.executeQuery ("SELECT * FROM CUJOSQL.WHERECUREX");

        while (rs.next ()) {
            System.out.println("Index " + rs.getInt(1) + " value " + rs.getString(2));
        }

        rs.close ();
        s.close();
        System.out.println("-----");
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}
}

```

Example: Change values with a statement through another statement's cursor: ➤ This is an example of how to change values with a statement through another statement's cursor.

Example: Change values with a statement through another statement's cursor

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;

public class UsingPositionedUpdate {
    public Connection connection = null;
    public static void main(java.lang.String[] args) {

        UsingPositionedUpdate test = new UsingPositionedUpdate();

        test.setup();
        test.displayTable();

        test.run();
        test.displayTable();

        test.cleanup();
    }

    /**
    Handle all the required setup work.
    */
    public void setup() {
        try {
            // Register the JDBC driver.
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

            connection = DriverManager.getConnection("jdbc:db2:*local");

            Statement s = connection.createStatement();
            try {
                s.executeUpdate("DROP TABLE CUJOSQL.WHERECUREX");
            } catch (SQLException e) {
                // Ignore problems here.
            }

            s.executeUpdate("CREATE TABLE CUJOSQL.WHERECUREX ( " +
                "COL_IND INT, COL_VALUE CHAR(20)) ");

            for (int i = 1; i <= 10; i++) {

```

```

        s.executeUpdate("INSERT INTO CUJOSQL.WHERECUREX VALUES(" + i + ", 'FIRST')");
    }

    s.close();

} catch (Exception e) {
    System.out.println("Caught exception: " + e.getMessage());
    e.printStackTrace();
}
}

/**
In this section, all the code to perform the testing should
be added. If only one connection to the database is required,
the global variable 'connection' can be used.
**/
public void run() {
    try {
        Statement stmt1 = connection.createStatement();

        // Update each value using next().
        stmt1.setCursorName("CUJO");
        ResultSet rs = stmt1.executeQuery ("SELECT * FROM CUJOSQL.WHERECUREX " +
            "FOR UPDATE OF COL_VALUE");

        System.out.println("Cursor name is " + rs.getCursorName());

        PreparedStatement stmt2 = connection.prepareStatement ("UPDATE "
            + " CUJOSQL.WHERECUREX
            SET COL_VALUE = 'CHANGED'
            WHERE CURRENT OF "
            + rs.getCursorName ());

        // Loop through the ResultSet and update every other entry.
        while (rs.next ()) {
            if (rs.next())
                stmt2.execute ();
        }

        // Clean up the resources after they have been used.
        rs.close ();
        stmt2.close ();

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

/**
In this section, put all clean-up work for testing.
**/
public void cleanup() {
    try {
        // Close the global connection opened in setup().
        connection.close();
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}

```

```

    }

/**
Display the contents of the table.
**/
public void displayTable()
{
    try {
        Statement s = connection.createStatement();
        ResultSet rs = s.executeQuery ("SELECT * FROM CUJOSQL.WHERECUREX");

        while (rs.next ()) {
            System.out.println("Index " + rs.getInt(1) + " value " + rs.getString(2));
        }

        rs.close ();
        s.close();
        System.out.println("-----");
    } catch (Exception e) {
        System.out.println("Caught exception: ");
        e.printStackTrace();
    }
}
}

```

Create ResultSets: ➤ To create a ResultSet object, you can use executeQuery methods from the Statement, PreparedStatement, or CallableStatement interfaces. There are, however, other available methods. For example, DatabaseMetaData methods such as getColumnNames, getTables, getUDTs, getPrimaryKeys, and so on, return ResultSets. It is also possible to have a single SQL statement return multiple ResultSets for processing. You can also use the getResultSet method to retrieve a ResultSet object after calling the execute method provided by the Statement, PreparedStatement, or CallableStatement interfaces.

See Example: Create a procedure with multiple ResultSets for more information.

Close ResultSets: While a ResultSet object is automatically closed when the Statement object with which it is associated closes, it is recommended that you close ResultSet objects when you are finished using them. By doing so, you immediately free internal database resources that can increase application throughput.

It is also important to close ResultSets generated by DatabaseMetaData calls. Because you do not directly have access to the Statement object that was used to create these ResultSets, you do not call close on the Statement object directly. These objects are linked together in such a way that the JDBC driver closes the internal Statement object when you close the external ResultSet object. When these objects are not closed manually, the system continues to work; however, it uses more resources than is necessary.

Note: The holdability characteristic of ResultSets can also close ResultSets automatically on your behalf. Calling close multiple times on a ResultSet object is allowed. ⏪

Example: ResultSet interface for IBM Developer Kit for Java: This is an example of how to use the ResultSet interface.

Example 1: ResultSet interface

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;

/**

```

ResultSetExample.java

This program demonstrates using a `ResultSetMetaData` and a `ResultSet` to display all the data in a table even though the program that gets the data does not know what the table is going to look like (the user passes in the values for the table and library).

```
/**/
public class ResultSetExample {

    public static void main(java.lang.String[] args)
    {
        if (args.length != 2) {
            System.out.println("Usage: java ResultSetExample <library> <table>");
            System.out.println(" where <library> is the library that contains <table>");
            System.exit(0);
        }

        Connection con = null;
        Statement s = null;
        ResultSet rs = null;
        ResultSetMetaData rsmd = null;

        try {
            // Get a database connection and prepare a statement.
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            con = DriverManager.getConnection("jdbc:db2:*local");

            s = con.createStatement();

            rs = s.executeQuery("SELECT * FROM " + args[0] + "." + args[1]);
            rsmd = rs.getMetaData();

            int colCount = rsmd.getColumnCount();
            int rowCount = 0;
            while (rs.next()) {
                rowCount++;
                System.out.println("Data for row " + rowCount);
                for (int i = 1; i <= colCount; i++)
                    System.out.println("  Row " + i + ": " + rs.getString(i));
            }

        } catch (Exception e) {
            // Handle any errors.
            System.out.println("Oops... we have an error... ");
            e.printStackTrace();
        } finally {
            // Ensure we always clean up. If the connection gets closed, the
            // statement under it closes as well.
            if (con != null) {
                try {
                    con.close();
                } catch (SQLException e) {
                    System.out.println("Critical error - cannot close connection object");
                }
            }
        }
    }
}
}
```

Example: ResultSetMetaData interface for IBM Developer Kit for Java: Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
```

```
/**
ResultSetMetaDataExample.java
```

This program demonstrates using a `ResultSetMetaData` and a `ResultSet` to display all the metadata about a `ResultSet` created querying a table. The user passes the value for the table and library in.

```

**/
public class ResultSetMetaDataExample {

    public static void main(java.lang.String[] args)
    {
        if (args.length != 2) {
            System.out.println("Usage: java ResultSetMetaDataExample <library> <table>");
            System.out.println("where <library> is the library that contains <table>");
            System.exit(0);
        }

        Connection con = null;
        Statement s = null;
        ResultSet rs = null;
        ResultSetMetaData rsmd = null;

        try {
            // Get a database connection and prepare a statement.
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
            con = DriverManager.getConnection("jdbc:db2:*local");

            s = con.createStatement();

            rs = s.executeQuery("SELECT * FROM " + args[0] + "." + args[1]);
            rsmd = rs.getMetaData();

            int colCount = rsmd.getColumnCount();
            int rowCount = 0;
            for (int i = 1; i <= colCount; i++) {
                System.out.println("Information about column " + i);
                System.out.println("  Name.....: " + rsmd.getColumnName(i));
                System.out.println("  Data Type.....: " + rsmd.getColumnType(i) +
                    " ( " + rsmd.getColumnTypeName(i) + " )");
                System.out.println("  Precision.....: " + rsmd.getPrecision(i));
                System.out.println("  Scale.....: " + rsmd.getScale(i));
                System.out.print ("  Allows Nulls..: ");
                if (rsmd.isNullable(i)==0)
                    System.out.println("false");
                else
                    System.out.println("true");
            }

        } catch (Exception e) {
            // Handle any errors.
            System.out.println("Oops... we have an error... ");
            e.printStackTrace();
        } finally {
            // Ensure we always clean up. If the connection gets closed, the
            // statement under it closes as well.
            if (con != null) {
                try {
                    con.close();
                } catch (SQLException e) {
                    System.out.println("Critical error - cannot close connection object");
                }
            }
        }
    }
}

```


JDBC object pooling

» Object pooling is the most common topic to come up when discussing Java^(TM) Database Connectivity (JDBC) and performance. Since many objects used in JDBC are expensive to create such as Connection, Statement, and ResultSet objects, significant performance benefits can be achieved by reusing these objects instead of creating every time you need them.

Many applications already handle object pooling on your behalf. For example, WebSphere has extensive support for pooling JDBC objects and allows you to control how the pool is managed. Because of this, you can get the functionality you want without being concerned about your own pooling mechanisms. However, when the support is not provided, you must find a solution for all but trivial applications.

To use object pooling in your JDBC programs, see the following:

Use DataSource support for object pooling

You can use DataSources to have multiple applications share a common configuration for accessing a database. This is accomplished by having each application reference the same DataSource name.

ConnectionPoolDataSource properties

You can configure the ConnectionPoolDataSource interface by using the set of properties that it provides.

DataSource-based statement pooling

You can use statement pooling within a connection pool. The UDBConnectionPoolDataSource interface's maxStatements property allows the DataSource to specify how many statements can be pooled under a connection.

Build your own pooling solution

You can develop your own connection and statement pooling without requiring support for DataSources or relying on another product.



Use DataSource support for object pooling: » Using DataSources allows you to have multiple applications share a common configuration for accessing the database. This is accomplished by having each application reference the same DataSource name.

By using DataSources, many applications can be changed from a central location. For example, if you change the name of a default library used by all your applications and you have used a single DataSource to obtain connections for all of them, you can update the name of the collection in that DataSource. All of your applications then start using the new default library.

When using DataSources to obtain connections for an application, you can use the native JDBC driver's built-in support for connection pooling. This support is provided as an implementation of the ConnectionPoolDataSource interface.

Pooling is accomplished by handing out "logical" Connection objects instead of physical Connection objects. A **logical Connection object** is a connection object that is returned by a pooled Connection object. Each logical connection object acts as a temporary handle to the physical connection represented by the pooled connection object. To the application, when the Connection object is returned, there is no noticeable difference between the two. The subtle difference comes when you call the close method on the Connection object. This call invalidates the logical connection and returns the physical connection to the pool where another application is able to use the physical connection. This technique lets many logical connection objects reuse a single physical connection.

Set up connection pooling: Connection pooling is accomplished by creating a DataSource object that references a ConnectionPoolDataSource object. ConnectionPoolDataSource objects have properties that can be set for handling various aspects of pool maintenance.

Refer to the example on how to set up connection pooling with UDBDataSource and UDBConnectionPoolDataSource more details. You can also see the Java Naming and Directory Interface (JNDI) for details about the role JNDI plays in this example.

From the example, the link that binds the two DataSource objects together is the dataSourceName. The link tells the DataSource object to defer establishing connections to the ConnectionPoolDataSource object that manages pooling automatically.

Pooling and non-pooling applications: There is no difference between an application that uses Connection pooling and one that does not. Therefore, pooling support can be added after the application code is complete, without making any changes to the application code.

Refer to Example: Test the performance of connection pooling for more details.

The following is output from running the previous program locally during development.

```
Start timing the non-pooling DataSource version...
Time spent: 6410
```

```
Start timing the pooling version...
Time spent: 282
```

```
Java program completed.
```

By default, a UDBConnectionPoolDataSource pools a single connection. If an application needs a connection several times and only needs one connection at a time, using UDBConnectionPoolDataSource is a perfect solution. If you need many simultaneous connections, you must configure your ConnectionPoolDataSource to match your needs and resources. <<

Example: Set up connection pooling with UDBDataSource and UDBConnectionPoolDataSource:

➤ This is an example of how to use connection pooling with UDBDataSource and UDBConnectionPoolDataSource.

Example: Set up connection pooling with UDBDataSource and UDBConnectionPoolDataSource

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import javax.naming.*;
import com.ibm.db2.jdbc.app.UDBDataSource;
import com.ibm.db2.jdbc.app.UDBConnectionPoolDataSource;

public class ConnectionPoolingSetup
{
    public static void main(java.lang.String[] args)
        throws Exception
    {
        // Create a ConnectionPoolDataSource implementation
        UDBConnectionPoolDataSource cpds = new UDBConnectionPoolDataSource();
        cpds.setDescription("Connection Pooling DataSource object");

        // Establish a JNDI context and bind the connection pool data source
        Context ctx = new InitialContext();
        ctx.rebind("ConnectionSupport", cpds);

        // Create a standard data source that references it.
```

```

        UDBDataSource ds = new UDBDataSource();
        ds.setDescription("DataSource supporting pooling");
        ds.setDataSourceName("ConnectionSupport");
        ctx.rebind("PoolingDataSource", ds);
    }
}

```



Example: Test the performance of connection pooling: ➤ This is an example of how to test the performance of the pooling example against the performance of the non-pooling example.

Example: Test the performance of connection pooling

Note: Read the Code example disclaimer for important legal information.

```

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

public class ConnectionPoolingTest
{
    public static void main(java.lang.String[] args)
    throws Exception
    {
        Context ctx = new InitialContext();
        // Do the work without a pool:
        DataSource ds = (DataSource) ctx.lookup("BaseDataSource");
        System.out.println("\nStart timing the non-pooling DataSource version...");

        long startTime = System.currentTimeMillis();
        for (int i = 0; i < 100; i++) {
            Connection c1 = ds.getConnection();
            c1.close();
        }
        long endTime = System.currentTimeMillis();
        System.out.println("Time spent: " + (endTime - startTime));

        // Do the work with pooling:
        ds = (DataSource) ctx.lookup("PoolingDataSource");
        System.out.println("\nStart timing the pooling version...");

        startTime = System.currentTimeMillis();
        for (int i = 0; i < 100; i++) {
            Connection c1 = ds.getConnection();
            c1.close();
        }
        endTime = System.currentTimeMillis();
        System.out.println("Time spent: " + (endTime - startTime));
    }
}


```



ConnectionPoolDataSource properties: ➤ The ConnectionPoolDataSource interface provides a set of properties for its configuration. Descriptions of these properties are provided in the following table.

Property	Description
initialPoolSize	When the pool is first instantiated, this property determines how many connections are placed into the pool. If this value is specified outside the range of minPoolSize and maxPoolSize, either minPoolSize or maxPoolSize is used as the number of initial connections to create.
maxPoolSize	<p>As the pool is used, more connections may be requested than the pool has in it. This property specifies the maximum number of connections allowed to be created in the pool.</p> <p>Applications do not "block" and wait for a connection to be returned to the pool when the pool is at its maximum size and all connections are in use. Instead, the JDBC driver constructs a new connection based on the DataSource properties and returns the connection.</p> <p>If a maxPoolSize of 0 is specified, the pool is allowed to grow unbounded as long as the system has resources available to hand out.</p>
minPoolSize	<p>Spikes in using the pool can cause it to increase the number of connections in it. If the activity level diminishes to the point where some Connections are never pulled out of the pool, the resources are being taken up for no particular reason.</p> <p>In such cases, the JDBC driver has the ability to release some of the connections that it has accumulated. This property allows you to tell the JDBC to release connections, ensuring that it always has a certain number of connections available to use.</p> <p>If a minPoolSize of 0 is specified, it is possible for the pool to free all of its connections and for the application to actually pay for the connection time for each connection request.</p>
maxIdleTime	<p>Connections keep track of how long they have been sitting around without being used. This property specifies how long an application allows connections to be unused before they are released (that is, there are more connections than are needed).</p> <p>This property is a time in seconds and does not specify when the actual close occurs. It specifies when enough time has passed that the connection should be released.</p>
propertyCycle	This property represents the number of seconds that are allowed to pass between the enforcement of these rules.

Note: Setting either the maxIdleTime or the propertyCycle time to 0 means that the JDBC driver does not check for connections to be removed from the pool on its own. The rules specified for initial, min, and max size are still enforced.

When maxIdleTime and propertyCycle are not 0, a management thread is used to watch over the pool. The thread wakes up every propertyCycle second and checks all the connections in the pool to see which ones have been there without being used for more than maxIdleTime seconds. Connections fitting this criterion are removed from the pool until the minPoolSize is reached. 

DataSource-based statement pooling: ➤ Another property that is available on the `UDBConnectionPoolDataSource` interface is `maxStatements`. This property allows for statement pooling within the connection pool. Statement pooling only has an effect on `PreparedStatement`s and `CallableStatement`s. Statement objects are not pooled.

The implementation of statement pooling is similar to that of connection pooling. When the application calls `Connection.prepareStatement("select * from tablex")`, the pooling module checks if the `Statement` object has already been prepared under the connection. If it has, a logical `PreparedStatement` object is handed to you instead of the physical object. When you call `close`, the `Connection` object is returned to the pool, the logical `Connection` object is thrown away, and the `Statement` object can be reused.

The `maxStatements` property allows the `DataSource` to specify how many statements can be pooled under a connection. A value of 0 indicates that statement pooling should not be used. When the statement pool is full, a least recently used algorithm is applied to determine which statement is to be thrown out.

Example: Test the performance of two `DataSource`s tests one `DataSource` that uses connection pooling only and the other `DataSource` that uses statement and connection pooling.

The following example is output from running this program locally during development.

```
Deploying statement pooling data source
Start timing the connection pooling only version...
Time spent: 26312
```

```
Starting timing the statement pooling version...
Time spent: 2292
Java program completed <<
```

Example: Test the performance of two `DataSource`s: ➤ This is an example of testing one `DataSource` that uses connection pooling only and the other `DataSource` that uses statement and connection pooling.

Example: Test the performance of two `DataSource`s

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import javax.naming.*;
import java.util.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.UDBDataSource;
import com.ibm.db2.jdbc.app.UDBConnectionPoolDataSource;

public class StatementPoolingTest
{

    public static void main(java.lang.String[] args)
    throws Exception
    {
        Context ctx = new InitialContext();

        System.out.println("deploying statement pooling data source");
        deployStatementPoolDataSource();

        // Do the work with connection pooling only.
        DataSource ds = (DataSource) ctx.lookup("PoolingDataSource");
        System.out.println("\nStart timing the connection pooling only version...");

        long startTime = System.currentTimeMillis();
        for (int i = 0; i < 100; i++) {
```

```

        Connection c1 = ds.getConnection();
        PreparedStatement ps = c1.prepareStatement("select * from qsys2.sysprocs");
        ResultSet rs = ps.executeQuery();
        c1.close();
    }
    long endTime = System.currentTimeMillis();
    System.out.println("Time spent: " + (endTime - startTime));

    // Do the work with statement pooling added.
    ds = (DataSource) ctx.lookup("StatementPoolingDataSource");
    System.out.println("\nStart timing the statement pooling version...");

    startTime = System.currentTimeMillis();
    for (int i = 0; i < 100; i++) {
        Connection c1 = ds.getConnection();
        PreparedStatement ps = c1.prepareStatement("select * from qsys2.sysprocs");
        ResultSet rs = ps.executeQuery();
        c1.close();
    }
    endTime = System.currentTimeMillis();
    System.out.println("Time spent: " + (endTime - startTime));
}

private static void deployStatementPoolDataSource()
throws Exception
{
    // Create a ConnectionPoolDataSource implementation
    UDBConnectionPoolDataSource cpds = new UDBConnectionPoolDataSource();
    cpds.setDescription("Connection Pooling DataSource object with Statement pooling");
    cpds.setMaxStatements(10);

    // Establish a JNDI context and bind the connection pool data source
    Context ctx = new InitialContext();
    ctx.rebind("StatementSupport", cpds);

    // Create a standard datasource that references it.
    UDBDataSource ds = new UDBDataSource();
    ds.setDescription("DataSource supporting statement pooling");
    ds.setDataSourceName("StatementSupport");
    ctx.rebind("StatementPoolingDataSource", ds);
}
}


```

Build your own connection pooling: ➤ You can develop your own connection and statement pooling without requiring support for DataSources or relying on another product.

The pooling techniques are demonstrated on a small Java application, but are equally applicable to servlets or large n-tiered applications. This example is used to demonstrate the performance issues.

The demonstration application has two functions:

- To insert a new index and name into a database table.
- To read the name for a given index from the table.

The complete code to the application can be downloaded from IBM's Developer Kit for Java JDBC Web Page  .

The example application does not perform well. Running 100 calls to the getValue method and 100 calls to the putValue method through this code took an average of 31.86 seconds on a standard workstation.

The problem is that there is too much database work for every request. That is, you get a connection, get a statement, process the statement, close the statement, and close the connection. Instead of discarding everything after each request, there must be a way to reuse portions of this process. **Connection pooling** is replacing the create connection code with code to obtain a connection from the pool, and then replacing the close connection code with code to return the connection to the pool for use.

The connection pool's constructor creates the connections and places them in the pool. The pool class has take and put methods for locating a connection to use and for returning the connection to the pool when done working with the connection. These methods are synchronized because the pool object is a shared resource, but you do not want multiple threads to simultaneously try to manipulate the pooled resources.

There is a change to the calling code for the `getValue` method. The `putValue` method is not shown, but the exact change is made to it and is available from IBM's Developer Kit for Java JDBC Web Page [👉](#). The instantiation of the connection pool object is also not shown. You can call the constructor and pass in the number of connection objects that you want in the pool. This step should be done when you start up the application.

Running the previous application (that is, having 100 `getValue` method and 100 `putValue` method requests) with these changes took an average of 13.43 seconds with the connection pooling code in place. The processing time for the workload is cut by more than half the original processing time without connection pooling.

Build your own statement pooling: When using connection pooling, time is wasted when creating and closing a statement when each statement is processed. This is another example of wasting an object that can be reused.

To reuse an object, you can use the prepared statement class. In most applications, the same SQL statements are reused with minor changes. For example, one iteration through an application might generate the following query:

```
SELECT * from employee where salary > 100000
```

The next iteration might generate the following query:

```
SELECT * from employee where salary > 50000
```

This is the same query, but it uses a different parameter. Both queries can be accomplished with the following query:

```
SELECT * from employee where salary > ?
```

You can then set the parameter marker (denoted by the question mark) to 100000 when processing the first query and 50000 when processing the second query. This enhances performance for three reasons beyond what the connection pool can offer:

- Fewer objects are created. A `PreparedStatement` object is created and reused instead of creating a `Statement` object for every request. Therefore, you run fewer constructors.
- The database work to set up the SQL statement (called the **prepare**) can be reused. Preparing SQL statements is reasonably expensive as it involves determining what the SQL statement text says and how the system should accomplish the task requested.
- When removing the additional object creations, there is a benefit that is not often considered. There is no need to destroy what was not created. This model is easier on the Java garbage collector and also benefits performance over time with many users.

The demonstration program can be changed to pool `PreparedStatement` objects instead of `Connections`. Changing the program allows you to reuse more object and improve performance. You can begin by writing the class that contains the objects to be pooled. This class must encapsulate the various resources

to be used. For the connection pool example, the Connection was the only pooled resource, so there was no need for an encapsulating class. Each pooled object must contain a Connection and two PreparedStatements. You can then create a pool class that contains database access objects instead of connections.

Finally, the application must change to obtain a database access object and specify which resource from the object it wants to use. Other than specifying the specific resource, the application remains the same.

With this change, the same test run now takes an average of 0.83 seconds. This time is about 38 times faster than the original version of the program.

Considerations: Performance improves through replication. If an item is not reused, then it is wasting resources to pool it.

Most applications contain critical sections of code. Typically, an application uses 80 to 90 percent of its processing time on only 10 to 20 percent of the code. If there are 10,000 SQL statements potentially used in an application, not all of them are pooled. The objective is to identify and pool the SQL statements that are used in the application's critical sections of code.

Creating objects in a Java implementation can carry a heavy cost. The pooling solution can be used with advantage. Objects used in the process are created at the beginning, before other users attempt to use the system. These objects are reused as often as required. Performance is excellent and it is possible to fine-tune the application over time to facilitate its use for greater numbers of users. As a result, more objects are pooled. Moreover, it permits more efficient multithreading of the application's database access to gain greater throughput.

Java (using JDBC) is based on dynamic SQL and tends to be slow. Pooling can minimize this problem. By preparing the statements at startup, access to the database can be rendered static. There is little difference in performance between dynamic and static SQL after the statement is prepared.

The performance of database access in Java can be efficient and can be accomplished without sacrificing object-oriented design or code maintainability. Writing code to build statement and connection pooling is not difficult. Furthermore, the code can be changed and enhanced to support multiple applications and application types (Web-based, client/server) and so on. <<

Batch updates

>> A new feature in JDBC 2.0 is batch update support. This feature allows any updates to the database to be passed as a single transaction between the user program and the database. This procedure can significantly improve performance when many updates must be performed at once. For example, if a large company requires its newly hired employees to start work on a Monday, this requirement makes it necessary to process many updates (in this case, insertions) to the employee database at one time. Creating a batch of updates and submitting them to the database as one unit can save you processing time.

There are two types of batch updates:

- Batch updates that use Statement objects.
- Batch updates that use PreparedStatement objects.

To use batch update support, see the following:

Statement batch update

Before performing a statement batch update, you must ensure that auto-commit is turned off. When the auto-commit setting is off, you can create a standard Statement object. You can then add the

statements to the batch with the `addBatch` method. Once you have added all the statements you want to the batch, you can process all of them with the `executeBatch` method or empty the batch at anytime with the `clearBatch` method.

PreparedStatement batch update

A `PreparedStatement` batch is similar to the `Statement` batch. However, a `PreparedStatement` batch always works off the same prepared statement, and you only change the parameters to that statement.


BatchUpdateException

When a call to the `executeBatch` method fails, a `BatchUpdateException` is thrown. The `BatchUpdateException` allows you to call all the same methods you have always called to receive the message, the `SQLState`, and vendor code. `BatchUpdateException` also provides the `getUpdateCounts` method that returns an integer array. The integer array contains update counts from all the statements in the batch that were processed up to the point where the failure occurred.

Blocked insert support

You can use a blocked insert as an `iSeries` operation to insert several rows into a database table at a time.



Statement batch update:  To perform a `Statement` batch update, you must turn off auto-commit. In Java^(TM) Database Connectivity (JDBC), auto-commit is on by default. Auto-commit means any updates to the database are committed after each SQL statement is processed. If you want to treat a group of statements being handed to the database as one functional group, you do not want the database committing each statement individually. If you do not turn off auto-commit and a statement in the middle of the batch fails, you cannot roll back the entire batch and try it again because half of the statements have been made final. Further, the additional work of committing each statement in a batch creates a lot of overhead. See `Transactions` for more details.


After turning off auto-commit, you can create a standard `Statement` object. Instead of processing statements with methods such as `executeUpdate`, you add them to the batch with the `addBatch` method. Once you have added all the statements you want to the batch, you can process all of them with the `executeBatch` method. You can empty the batch at anytime with the `clearBatch` method.

The following example shows how you can use these methods:

Example: Statement batch update

Note: Read the Code example disclaimer for important legal information.

```
connection.setAutoCommit(false);
Statement statement = connection.createStatement();
statement.addBatch("INSERT INTO TABLEX VALUES(1, 'Cujo')");
statement.addBatch("INSERT INTO TABLEX VALUES(2, 'Fred')");
statement.addBatch("INSERT INTO TABLEX VALUES(3, 'Mark')");
int [] counts = statement.executeBatch();
connection.commit();
```

In this example, an array of integers is returned from the `executeBatch` method. This array has one integer value for each statement that is processed in the batch. If values are being inserted into the database, the value for each statement is 1 (that is, assuming successful processing). However, some of the statements may be update statements that affect multiple rows. If you put any statements in the batch other than `INSERT`, `UPDATE`, or `DELETE`, an exception occurs. 

PreparedStatement batch update: ➤ A preparedStatement batch is similar to the Statement batch; however, a preparedStatement batch always works off the same "prepared" statement, and you only change the parameters to that statement. The following is an example that uses a preparedStatement batch.

Example: PreparedStatement batch update

Note: Read the Code example disclaimer for important legal information.

```
connection.setAutoCommit(false);
PreparedStatement statement =
    connection.prepareStatement("INSERT INTO TABLEX VALUES(?, ?)");
statement.setInt(1, 1);
statement.setString(2, "Cujo");
statement.addBatch();
statement.setInt(1, 2);
statement.setString(2, "Fred");
statement.addBatch();
statement.setInt(1, 3);
statement.setString(2, "Mark");
statement.addBatch();
int [] counts = statement.executeBatch();
connection.commit();
```



BatchUpdateException: ➤ An important consideration of batch updates is what action to take when a call to the executeBatch method fails. In this case, a new type of exception, called BatchUpdateException, is thrown. The BatchUpdateException is a subclass of SQLException and it allows you to call all the same methods you have always called to receive the message, the SQLState, and vendor code. BatchUpdateException also provides the getUpdateCounts method that returns an integer array. The integer array contains update counts from all the statements in the batch that were processed up to the point where the failure occurred. The array length tells you which statement in the batch failed. For example, if the array returned in the exception has a length of three, the fourth statement in the batch failed. Therefore, from the single BatchUpdateException object that is returned, you can determine the update counts for all the statements that were successful, which statement failed, and all the information about the failure.

Currently, the standard performance of processing batched updates is equivalent to the performance of processing each statement independently. You can refer to Blocked insert support for more information on optimized support for batch updates. You should still use the new model when coding and take advantage of future performance optimizations.

Note: In the JDBC 2.1 specification, a different option is provided for how exception conditions for batch updates are handled. JDBC 2.1 introduces a model where the processing batch continues after a batch entry fails. A special update count is placed in the array of update count integers that is returned for each entry that fails. This allows large batches to continue processing even though one of their entries fails. See the JDBC 2.1 or JDBC 3.0 specification for details on these two modes of operation. By default, the native JDBC driver uses the JDBC 2.0 definition. The driver provides a Connection property that is used when using DriverManager to establish connections. The driver also provides a DataSource property that is used when using DataSources to establish connections. These properties allow applications to choose how they want batch operations to handle failures. ⬅

Blocked insert support: ➤ A **blocked insert** is a special type of operation on an iSeries server that provides a highly optimized way to insert several rows into a database table at a time. Blocked inserts can be thought of as a subset of batched updates. Batched updates can be any form of an update request, but blocked inserts are specific. However, blocked insert types of batched updates are common; the native JDBC driver has been changed to take advantage of this feature.

Because of system restrictions when using blocked insert support, the default setting for the native JDBC driver is to have blocked insert disabled. It can be enabled through a Connection property or a DataSource property. Most of the restrictions when using a blocked insert can be checked and handled on your behalf, but a few restrictions cannot; thus, this is the reason for turning off blocked insert support by default. The list of restrictions is as follows:

- The SQL statement used must be an INSERT statement with a VALUES clause, meaning that it is not an INSERT statement with SUBSELECT. The JDBC driver recognizes this restriction and takes the appropriate course of action.
- A PreparedStatement must be used, meaning that there is no optimized support for Statement objects. The JDBC driver recognizes this restriction and takes the appropriate course of action.
- The SQL statement must specify parameter markers for all the columns in the table. This means that you cannot either use constant values for a column or allow the database to insert default values for any of the columns. The JDBC driver does not have a mechanism to handle testing for specific parameter markers in your SQL statement. If you set the property to perform optimized blocked insertions and you do not avoid defaults or constants in your SQL statements, the values that end up in the database table are not correct.
- The connection must be to the local system. This means that a connection using DRDA to access a remote system cannot be used because DRDA does not support a blocked insert operation. The JDBC driver does not have a mechanism to handle testing for a connection to a local system. If you set the property to perform an optimized blocked insertion and you attempt to connect to a remote system, the processing of the batch update fails.

This code example shows how to enable support for blocked insert processing. The only difference between this code and a version that would not use blocked insert support is use `block insert=true` that is added to the Connection URL.

Example: Blocked insert processing

Note: Read the Code example disclaimer for important legal information.

```
// Create a database connection
Connection c = DriverManager.getConnection("jdbc:db2:*local;use block insert=true");
BigDecimal bd = new BigDecimal("123456");

// Create a PreparedStatement to insert into a table with 4 columns
PreparedStatement ps =
    c.prepareStatement("insert into cujosql.xxx values(?, ?, ?, ?)");

// Start timing...
for (int i = 1; i <= 10000; i++) {
    ps.setInt(1, i);           // Set all the parameters for a row
    ps.setBigDecimal(2, bd);
    ps.setBigDecimal(3, bd);
    ps.setBigDecimal(4, bd);
    ps.addBatch();           //Add the parameters to the batch
}

// Process the batch
int[] counts = ps.executeBatch();

// End timing...
```

In similar test cases, a blocked insert is several times faster than performing the same operations when a blocked insert is not used. For example, the test performed on the previous code was nine time faster using blocked inserts. Cases that only use primitive types instead of objects can be up to sixteen times faster. In applications where there is a significant amount of work going on, expectations should be appropriately changed. <<

Advanced data types

➤ There are several new data types called SQL3 data types that are provided in the iSeries database with the V4R4 e-PACK. Java^(TM) Database Connectivity (JDBC) 2.0 and higher provide support for working with these data types that are a part of the SQL99 standard.

The SQL3 data types give you a tremendous amount of flexibility. They are ideal for storing serialized Java objects, Extensible Markup Language (XML) documents, and multimedia data such as songs, product pictures, employee photographs, and movie clips.

Distinct types: The **distinct type** is a user-defined type that is based on a standard database type. For example, you can define a Social Security Number type, SSN, that is a CHAR(9) internally. The following SQL statement creates such a DISTINCT type.

```
CREATE DISTINCT TYPE CUJOSQL.SSN AS CHAR(9)
```

A distinct type always maps to a built-in data type. For more information on how and when to use distinct types in the context of SQL, consult the SQL reference manuals.

To use distinct types in JDBC, you access them the same way that you access an underlying type. The `getUDTs` method is a new method that allows you to query what distinct types are available on the system. This example program shows the following:

- The creation of a distinct type.
- The creation of a table that uses it.
- The use of a `PreparedStatement` to set a distinct type parameter.
- The use of a `ResultSet` to return a distinct type.
- The use of the metadata Application Programming Interface (API) call to `getUDTs` to learn about a distinct type.

Large Objects: There are three types of Large Objects (LOBs):

- Binary Large Objects (BLOBs)
- Character Large Objects (CLOBs)
- Double Byte Character Large Objects (DBCLOBs)

DBCLOBs are similar to CLOBs except for their internal storage representation of the character data. Because Java and JDBC externalize all character data as Unicode, there is only support in JDBC for CLOBs. DBCLOBs work interchangeable with the CLOB support from a JDBC perspective.

Binary Large Objects: In many ways, a Binary Large Object (BLOB) column is similar to a CHAR FOR BIT DATA column that can be made large. You can store anything in these columns that can be represented as a stream of nontranslated bytes. Often, BLOB columns are used to store serialized Java objects, pictures, songs, and other binary data.

You can use BLOBs the same way you would use other standard database types. You can pass them to stored procedures, use them in prepared statements, and update them in result sets. The `PreparedStatement` class has a `setBlob` method for passing BLOBs to the database, and the `ResultSet` class adds a `getBlob` class for retrieving them from the database. A BLOB is represented in a Java program by a BLOB object that is a JDBC interface.

Refer to `Write code that uses BLOBs` for more information on how to use BLOBs.

Character Large Objects: Character Large Objects (CLOBs) are the character data complement to BLOBs. Instead of storing data in the database without translation, the data is stored in the database as text and is processed the same way as a CHAR column. As with BLOBs, JDBC 2.0 provides functions for dealing directly with CLOBs. The `PreparedStatement` interface contains a `setClob` method and the `ResultSet` interface contains a `getClob` method.

Refer to Write code that uses CLOBs for more information on how to use CLOBs.

Although BLOB and CLOB columns work like CHAR FOR BIT DATA and CHAR columns, this is conceptually how they work from an external user's perspective. Internally, they are different; because of the potentially enormous size of Large Object (LOB) columns, you typically work indirectly with data. For example, when a block of rows is fetched from the database, you do not move a block of LOBs to the ResultSet. You move pointers called LOB locators (that is, four-byte integers) into the ResultSet instead. However, it is not necessary to know about locators when working with LOBs in JDBC.

Datalinks: **Datalinks** are encapsulated values that contain a logical reference from the database to a file stored outside the database. Datalinks are represented and used from a JDBC perspective in two different ways, depending on whether you are using JDBC 2.0 or earlier, or you are using JDBC 3.0 or later.

Refer to Write code that uses Datalinks for more information on how to use Datalinks.

Unsupported SQL3 data types: There are other SQL3 data types that have been defined and for which the JDBC API provides support. These are ARRAY, REF, and STRUCT. Presently, iSeries servers do not support these types. Therefore, the JDBC driver does not provide any form of support for them. <<

Write code that uses BLOBs: >> There are a number of tasks that can be accomplished with database Binary Large Object (BLOB) columns through the Java^(TM) Database Connectivity (JDBC) Application Programming Interface (API). The following topics briefly discuss these tasks and include examples on how to accomplish them.

Read BLOBs from the database and insert BLOBs into the database: With the JDBC API, there are ways to get BLOBs out of the database and ways to put BLOBs into the database. However, there is no standardized way to create a Blob object. This is not a problem if your database is already full of BLOBs, but it poses a problem if you want to work with BLOBs from scratch through JDBC. Instead of defining a constructor for the Blob and Clob interfaces of the JDBC API, support is provided for placing BLOBs into the database and getting them out of the database directly as other types. For example, the setBinaryStream method can work with a database column of type Blob. This example shows some of the common ways that a BLOB can be put into the database or retrieved from the database.

Work with the Blob object API: BLOBs are defined in JDBC as an interface of which the various drivers provide implementations. This interface has a series of methods that can be used to interact with the Blob object. This example shows some of the common tasks that can be performed using this API. Consult the JDBC Javadoc for a complete list of available methods on the Blob object.

Use JDBC 3.0 support to update BLOBs: In JDBC 3.0, there is support for making changes to LOB objects. These changes can be stored into BLOB columns in the database. This example shows some of the tasks that can be performed with BLOB support in JDBC 3.0. <<

Example: BLOB: >> This is an example of how a BLOB can be put into the database or retrieved from the database.

Example: BLOB

Note: Read the Code example disclaimer for important legal information.

```
//////////////////////////////////////////
// PutGetBlobs is an example application
// that shows how to work with the JDBC
// API to obtain and put BLOBs to and from
// database columns.
//
// The results of running this program
// are that there are two BLOB values
// in a new table. Both are identical
```

```

// and contain 500k of random byte
// data.
////////////////////////////////////
import java.sql.*;
import java.util.Random;

public class PutGetBlobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        // Establish a Connection and Statement with which to work.
        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        // Clean up any previous run of this application.
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.BLOBTABLE");
        } catch (SQLException e) {
            // Ignore it - assume the table did not exist.
        }

        // Create a table with a BLOB column. The default BLOB column
        // size is 1 MB.
        s.executeUpdate("CREATE TABLE CUJOSQL.BLOBTABLE (COL1 BLOB)");

        // Create a PreparedStatement object that allows you to put
        // a new Blob object into the database.
        PreparedStatement ps = c.prepareStatement("INSERT INTO CUJOSQL.BLOBTABLE VALUES(?)");

        // Create a big BLOB value...
        Random random = new Random ();
        byte [] inByteArray = new byte[500000];
        random.nextBytes (inByteArray);

        // Set the PreparedStatement parameter. Note: This is not
        // portable to all JDBC drivers. JDBC drivers do not have
        // support when using setBytes for BLOB columns. This is used to
        // allow you to generate new BLOBs. It also allows JDBC 1.0
        // drivers to work with columns containing BLOB data.
        ps.setBytes(1, inByteArray);

        // Process the statement, inserting the BLOB into the database.
        ps.executeUpdate();

        // Process a query and obtain the BLOB that was just inserted out
        // of the database as a Blob object.
        ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.BLOBTABLE");
        rs.next();
        Blob blob = rs.getBlob(1);

        // Put that Blob back into the database through
        // the PreparedStatement.
        ps.setBlob(1, blob);
        ps.execute();

        c.close(); // Connection close also closes stmt and rs.
    }
}

```



Example: Update BLOBs:  This is an example of how to update BLOBs in your applications.

Example: Update BLOBs

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
// UpdateBlobs is an example application
// that shows some of the APIs providing
// support for changing Blob objects
// and reflecting those changes to the
// database.
//
// This program must be run after
// the PutGetBlobs program has completed.
////////////////////////////////////
import java.sql.*;

public class UpdateBlobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.BLOBTABLE");

        rs.next();
        Blob blob1 = rs.getBlob(1);
        rs.next();
        Blob blob2 = rs.getBlob(1);

        // Truncate a BLOB.
        blob1.truncate((long) 150000);
        System.out.println("Blob1's new length is " + blob1.length());

        // Update part of the BLOB with a new byte array.
        // The following code obtains the bytes that are at
        // positions 4000-4500 and set them to positions 500-1000.

        // Obtain part of the BLOB as a byte array.
        byte[] bytes = blob1.getBytes(4000L, 4500);

        int bytesWritten = blob2.setBytes(500L, bytes);

        System.out.println("Bytes written is " + bytesWritten);

        // The bytes are now found at position 500 in blob2
        long startInBlob2 = blob2.position(bytes, 1);

        System.out.println("pattern found starting at position " + startInBlob2);

        c.close(); // Connection close also closes stmt and rs.
    }
}
```

Example: Use BLOBs: ➤ This is an example of how to use BLOBs in your applications.

Example: Use BLOBs

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
// UseBlobs is an example application
// that shows some of the APIs associated
// with Blob objects.
//
// This program must be run after
// the PutGetBlobs program has completed.
////////////////////////////////////
import java.sql.*;

public class UseBlobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.BLOBTABLE");

        rs.next();
        Blob blob1 = rs.getBlob(1);
        rs.next();
        Blob blob2 = rs.getBlob(1);

        // Determine the length of a LOB.
        long end = blob1.length();
        System.out.println("Blob1 length is " + blob1.length());

        // When working with LOBs, all indexing that is related to them
        // is 1-based, and is not 0-based like strings and arrays.
        long startingPoint = 450;
        long endingPoint = 500;

        // Obtain part of the BLOB as a byte array.
        byte[] outByteArray = blob1.getBytes(startingPoint, (int)endingPoint);

        // Find where a sub-BLOB or byte array is first found within a
        // BLOB. The setup for this program placed two identical copies of
        // a random BLOB into the database. Thus, the start position of the
        // byte array extracted from blob1 can be found in the starting
        // position in blob2. The exception would be if there were 50
        // identical random bytes in the LOBs previously.
        long startInBlob2 = blob2.position(outByteArray, 1);

        System.out.println("pattern found starting at position " + startInBlob2);

        c.close(); // Connection close closes stmt and rs too.
    }
}
```


Write code that uses CLOBs: » There are a number of tasks that can be performed with database CLOB and DBCLOB columns through the Java[™] Database Connectivity (JDBC) Application Programming Interface (API). The following topics briefly discuss these tasks and include examples on how to accomplish them.

Read CLOBs from the database and insert CLOBs into the database: With the JDBC API, there are ways to get CLOBs out of the database and ways to put CLOBs into the database. However, there is no standardized way to create a Clob object. This is not a problem if your database is already full of CLOBs, but it poses a problem if you want to work with CLOBs from scratch through JDBC. Instead of defining a constructor for the Clob and Blob interfaces of the JDBC API, support is provided for placing CLOBs into the database and getting them out of the database directly as other types. For example, the `setCharacterStream` method can work with a database column of type Clob. This example shows some of the common ways that a CLOB can be put into the database or retrieved from the database.

Work with the Clob object API: CLOBs are defined in JDBC as an interface of which the various drivers provide implementations. This interface has a series of methods that can be used to interact with the Clob object. This example shows some of the common tasks that can be performed using this API. Consult the JDBC Javadoc for a complete list of available methods on the Clob object.

Use JDBC 3.0 support to update CLOBs: In JDBC 3.0, there is support for making changes to LOB objects. These changes can be stored into CLOB columns in the database. This example shows some of the tasks that can be performed with CLOB support in JDBC 3.0. «

Example: CLOB: » This is an example of how a CLOB can be put into the database or retrieved from the database.

Example: CLOB

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
// PutGetClobs is an example application
// that shows how to work with the JDBC
// API to obtain and put CLOBs to and from
// database columns.
//
// The results of running this program
// are that there are two CLOB values
// in a new table. Both are identical
// and contain about 500k of repeating
// text data.
////////////////////////////////////
import java.sql.*;

public class PutGetClobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        // Establish a Connection and Statement with which to work.
        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        // Clean up any previous run of this application.
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.CLOBTABLE");
        }
    }
}
```

```

    } catch (SQLException e) {
        // Ignore it - assume the table did not exist.
    }

    // Create a table with a CLOB column. The default CLOB column
    // size is 1 MB.
    s.executeUpdate("CREATE TABLE CUJOSQL.CLOBTABLE (COL1 CLOB)");

    // Create a PreparedStatement object that allow you to put
    // a new Clob object into the database.
    PreparedStatement ps = c.prepareStatement("INSERT INTO CUJOSQL.CLOBTABLE VALUES(?)");

    // Create a big CLOB value...
    StringBuffer buffer = new StringBuffer(500000);
    while (buffer.length() < 500000) {
        buffer.append("All work and no play makes Cujo a dull boy.");
    }
    String clobValue = buffer.toString();

    // Set the PreparedStatement parameter. This is not
    // portable to all JDBC drivers. JDBC drivers do not have
    // to support setBytes for CLOB columns. This is done to
    // allow you to generate new CLOBs. It also
    // allows JDBC 1.0 drivers a way to work with columns containing
    // Clob data.
    ps.setString(1, clobValue);

    // Process the statement, inserting the clob into the database.
    ps.executeUpdate();

    // Process a query and get the CLOB that was just inserted out of the
    // database as a Clob object.
    ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.CLOBTABLE");
    rs.next();
    Clob clob = rs.getClob(1);

    // Put that Clob back into the database through
    // the PreparedStatement.
    ps.setClob(1, clob);
    ps.execute();

    c.close(); // Connection close also closes stmt and rs.
}
}

```



Example: Update CLOBs:  This is an example of how to update CLOBs in your applications.

Example: Update CLOBs

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////
// UpdateClobs is an example application
// that shows some of the APIs providing
// support for changing Clob objects
// and reflecting those changes to the
// database.
//
// This program must be run after
// the PutGetClobs program has completed.
////////////////////////////////////
import java.sql.*;

```

```

public class UpdateClobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.CLOBTABLE");

        rs.next();
        Clob clob1 = rs.getClob(1);
        rs.next();
        Clob clob2 = rs.getClob(1);

        // Truncate a CLOB.
        clob1.truncate((long) 150000);
        System.out.println("Clob1's new length is " + clob1.length());

        // Update a portion of the CLOB with a new String value.
        String value = "Some new data for once";
        int charsWritten = clob2.setString(500L, value);

        System.out.println("Characters written is " + charsWritten);

        // The bytes can be found at position 500 in clob2
        long startInClob2 = clob2.position(value, 1);

        System.out.println("pattern found starting at position " + startInClob2);

        c.close(); // Connection close also closes stmt and rs.
    }
}

```

Example: Use CLOBs: ➤ This is an example of how to use CLOBs in your applications.

Example: Use CLOBs

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////
// UpdateClobs is an example application
// that shows some of the APIs providing
// support for changing Clob objects
// and reflecting those changes to the
// database.
//
// This program must be run after
// the PutGetClobs program has completed.
////////////////////////////////////
import java.sql.*;

public class UseClobs {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");

```

```

    } catch (Exception e) {
        System.exit(1); // Setup error.
    }

    Connection c = DriverManager.getConnection("jdbc:db2:*local");
    Statement s = c.createStatement();

    ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.CLOBTABLE");

    rs.next();
    Clob clob1 = rs.getClob(1);
    rs.next();
    Clob clob2 = rs.getClob(1);

    // Determine the length of a LOB.
    long end = clob1.length();
    System.out.println("Clob1 length is " + clob1.length());

    // When working with LOBs, all indexing that is related to them
    // is 1-based, and not 0-based like strings and arrays.
    long startingPoint = 450;
    long endingPoint = 50;

    // Obtain part of the CLOB as a byte array.
    String outString = clob1.getSubString(startingPoint, (int)endingPoint);
    System.out.println("Clob substring is " + outString);

    // Find where a sub-CLOB or string is first found within a
    // CLOB. The setup for this program placed two identical copies of
    // a repeating CLOB into the database. Thus, the start position of the
    // string extracted from clob1 can be found in the starting
    // position in clob2 if the search begins close to the position where
    // the string starts.
    long startInClob2 = clob2.position(outString, 440);

    System.out.println("pattern found starting at position " + startInClob2);

    c.close(); // Connection close also closes stmt and rs.
}
}

```

Write code that uses Datalinks: ➤ How you are going to work with Datalinks is dependent on what release you are working with. In JDBC 3.0, there is support to work directly with Datalink columns using the `getURL` and `putURL` methods. With previous JDBC versions, you had to work with Datalink columns as if they were String columns. Presently, the database does not support automatic conversions between Datalink and character data types. As a result, you have to perform some type casting in your SQL statements.

This example shows some of the basic tasks of working with Datalink columns. ⏪

Example: Datalink: ➤ This is an example of how to use datalinks in your applications.

Example: Datalink

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////
// PutGetDatalinks is an example application
// that shows how to use the JDBC
// API to handle datalink database columns.
////////////////////////////////////
import java.sql.*;
import java.net.URL;

```

```

import java.net.MalformedURLException;

public class PutGetDatalinks {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        // Establish a Connection and Statement with which to work.
        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        // Clean up any previous run of this application.
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.DLTABLE");
        } catch (SQLException e) {
            // Ignore it - assume the table did not exist.
        }

        // Create a table with a datalink column.
        s.executeUpdate("CREATE TABLE CUJOSQL.DLTABLE (COL1 DATALINK)");

        // Create a PreparedStatement object that allows you to add
        // a new datalink into the database. Since conversing
        // to a datalink cannot be accomplished directly in the database, you
        // can code the SQL statement to perform the explicit conversion.
        PreparedStatement ps = c.prepareStatement("INSERT INTO CUJOSQL.DLTABLE
            VALUES(DLVALUE( CAST(? AS VARCHAR(100))))");

        // Set the datalink. This URL points you to an article about
        // the new features of JDBC 3.0.
        ps.setString (1, "http://www-106.ibm.com/developerworks/java/library/j-jdbcnew/index.html");

        // Process the statement, inserting the CLOB into the database.
        ps.executeUpdate();

        // Process a query and obtain the CLOB that was just inserted out of the
        // database as a Clob object.
        ResultSet rs = s.executeQuery("SELECT * FROM CUJOSQL.DLTABLE");
        rs.next();
        String datalink = rs.getString(1);

        // Put that datalink value into the database through
        // the PreparedStatement. Note: This function requires JDBC 3.0
        // support.
        /*
        try {
            URL url = new URL(datalink);
            ps.setURL(1, url);
            ps.execute();
        } catch (MalformedURLException mue) {
            // Handle this issue here.
        }
        */

        rs = s.executeQuery("SELECT * FROM CUJOSQL.DLTABLE");
        rs.next();
        URL url = rs.getURL(1);
        System.out.println("URL value is " + url);
        */
    }
}

```

```

        c.close(); // Connection close also closes stmt and rs.
    }
}

```



Example: Distinct types: ➤ This is an example of how to use distinct types.

Example: Distinct types

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////
// This example program shows examples of
// various common tasks that can be done
// with distinct types.
////////////////////////////////////
import java.sql.*;

public class Distinct {
    public static void main(String[] args)
        throws SQLException
    {
        // Register the native JDBC driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (Exception e) {
            System.exit(1); // Setup error.
        }

        Connection c = DriverManager.getConnection("jdbc:db2:*local");
        Statement s = c.createStatement();

        // Clean up any old runs.
        try {
            s.executeUpdate("DROP TABLE CUJOSQL.SERIALNOS");
        } catch (SQLException e) {
            // Ignore it and assume the table did not exist.
        }

        try {
            s.executeUpdate("DROP DISTINCT TYPE CUJOSQL.SSN");
        } catch (SQLException e) {
            // Ignore it and assume the table did not exist.
        }

        // Create the type, create the table, and insert a value.
        s.executeUpdate("CREATE DISTINCT TYPE CUJOSQL.SSN AS CHAR(9)");
        s.executeUpdate("CREATE TABLE CUJOSQL.SERIALNOS (COL1 CUJOSQL.SSN)");

        PreparedStatement ps = c.prepareStatement("INSERT INTO CUJOSQL.SERIALNOS VALUES(?)");
        ps.setString(1, "399924563");
        ps.executeUpdate();
        ps.close();

        // You can obtain details about the types available with new metadata in
        // JDBC 2.0
        DatabaseMetaData dmd = c.getMetaData();

        int types[] = new int[1];
        types[0] = java.sql.Types.DISTINCT;

        ResultSet rs = dmd.getUDTs(null, "CUJOSQL", "SSN", types);
        rs.next();
        System.out.println("Type name " + rs.getString(3) +

```

```

        " has type " + rs.getString(4));

    // Access the data you have inserted.
    rs = s.executeQuery("SELECT COL1 FROM CUJOSQL.SERIALNOS");
    rs.next();
    System.out.println("The SSN is " + rs.getString(1));

    c.close(); // Connection close also closes stmt and rs.
}
}

```



RowSets

➤ RowSets were originally added to the Java^(TM) Database Connectivity (JDBC) 2.0 Optional Package. Unlike some of the better-known interfaces of the JDBC specification, the RowSet specification is designed to be more of a framework than an actual implementation. The RowSet interfaces define a set of core functionality that all RowSets have. RowSet implementation providers have considerable freedom to define the functionality that is needed to fit their needs in a specific problem space.

To implement Rowsets using the native JDBC driver, see the following:

RowSet characteristics

You can request certain properties to be satisfied by the RowSets. Common properties include the set of interfaces to be supported by the resulting rowset.

DB2JdbcRowSet

DB2JdbcRowSet is a connected RowSet that serves as a wrapper on a DB2ResultSet and provides event-handling support.

DB2CachedRowSet

DB2CachedRowSet is a disconnected RowSet that allows the DB2ResultSet data to be stored inside the object. Once the data is within the object, the underlying DB2Connection object can be closed and the DB2CachedRowSet can continue to be used. Find the following information regarding DB2CachedRowSet:

- Use DB2CachedRowSets
- Create and populate a DB2CachedRowSet
- Access DB2CachedRowSet data and cursor manipulation
- Change DB2CachedRowSet data and reflect changes back to the data source
- Other DB2CachedRowSet features



RowSet characteristics: ➤ You can request certain properties to be satisfied by the rowsets. Common properties include the set of interfaces to be supported by the resulting rowset.

RowSets are ResultSets: The RowSet interface extends the ResultSet interface which means that RowSets have the ability to perform all the functions that ResultSets can do. For example, RowSets can be scrollable and updateable.

RowSets can be disconnected from the database: There are two categories of RowSets:

- **Connected**

While connected RowSets are populated with data, they always have internal connections to the underlying database open and serve as wrappers around a ResultSet implementation.

- **Disconnected**

Disconnected RowSets are not required to maintain connections to their data source at all times. Disconnected RowSets can be detached from the database, be used in a variety of ways, and then be reconnected to the database to mirror any changes made to them.

RowSets are JavaBeans components: RowSets have support for event handling based on the JavaBeans event-handling model. They also have properties that can be set. These properties can be used by the RowSet to perform the following:

- Establish a connection to the database.
- Process an SQL statement.
- Determine features of the data that the RowSet represents and handle other internal features of the RowSet object.

RowSets are serializable: RowSets can be serialized and deserialized to allow them to flow over a network connection, be written out to a flat file (that is, a text document without any word processing or other structure characters), and so on. <<

DB2CachedRowSet: >> The DB2CachedRowSet object is a disconnected RowSet, meaning that it can be used without being connected to the database. Its implementation adheres closely to the description of a CachedRowSet.

The DB2CachedRowSet is a container for rows of data from a ResultSet. The DB2CachedRowSet holds all its own data so it does not need to maintain a connection to the database other than explicitly while reading or writing data to the database.

Use DB2CachedRowSets

You can use methods provided by DB2CachedRowSet to improve your database's performance by allowing several people to use the same data. You can also hand out common ResultSets to clients by creating a copy of table data that does not change.

Create and populate a DB2CachedRowSet

Find out how to create and place data into a DB2CachedRowSet by following these tasks:

- Use the populate method
- Use DB2CachedRowSet properties and DataSources
- Use DB2CachedRowSet properties and JDBC URLs
- Use the setConnection(Connection) method to use an existing database connection
- Use the execute(Connection) method to use an existing database connection
- Use the execute(int) method to group database requests

Access DB2CachedRowSet data and cursor manipulation

RowSets depend on ResultSet methods. For many operations such as DB2CachedRowSet data access and cursor movement, there is no difference at the application level between using a ResultSet and using a RowSet.

Change DB2CachedRowSet data and reflect changes back to the data source


The DB2CachedRowSet uses the same methods as the standard ResultSet interface for making changes to the data in the RowSet object. The DB2CachedRowSet provides the acceptChanges method that is used to reflect changes to the RowSet back to the database where the data came from.

Other DB2CachedRowSet features

The DB2CachedRowSet class has some additional functionality that makes it more flexible to use. With the methods provided by DB2CachedRowSet, you can perform the following tasks:

- Obtain collections from DB2CachedRowSets
- Create copies of RowSets
- Create shares for RowSets



Use DB2CachedRowSet:  Because the DB2CachedRowSet object can be disconnected and serialized, it is useful in environments where it is not always practical to run a full JDBC driver (for example, on Personal Digital Assistants (PDAs) and Java^(TM)-enabled cell phones).

Since the DB2CachedRowSet object is contained in memory and its data is always known, it can serve as a highly optimized form of a scrollable ResultSet for applications. Whereas DB2ResultSets that are scrollable typically pay a performance penalty because their random movements interfere with the JDBC driver's ability to cache rows of data, RowSets do not have this issue.

Two methods are provided on DB2CachedRowSet that create new RowSets:

- The createCopy method creates a new RowSet that is identical to the copied one.
- The createShared method creates a new RowSet that shares the same underlying data as the original.

You can use the createCopy method to hand out common ResultSets to clients. If the table data is not changing, creating a copy of a RowSet and passing it to each client is more efficient than running a query against the database each time.

You can use the createShared method to improve your database's performance by allowing several people to use the same data. For example, assume that you have a Web site that shows the top twenty best-selling products on your home page when a customer connects. You want the information on your main page to be updated regularly, but running the query to get the most frequently purchased items every time a customer visits your main page is not practical. Using the createShared method, you can effectively create "cursors" for each customer without having to either process the query again or store an enormous amount of information in memory. When appropriate, the query to find the most frequently purchased products can be run again. The new data can populate the RowSet that is used to create the shared cursors and the servlets can use them.

DB2CachedRowSets provide a delayed processing feature. This feature allows multiple query requests to be grouped together and processed against the database as a single request. This is an example of using DB2CachedRowSets to eliminate some of the computational stress that the database would otherwise be under.

Because the RowSet must keep careful track of any changes that happen to it so that they are reflected back to the database, there is support for functions that undo changes or allow you to see all changes have been made. For example, there is a showDeleted method that can be used to tell the RowSet to let you fetch deleted rows. There are also cancelRowInsert and cancelRowDelete methods to undo row insertions and deletions, respectfully, after they have been made.

The DB2CachedRowSet object offers better interoperability with other Java APIs because of its event handling support and its toCollection methods that allow a RowSet or a portion of it to be converted into a Java collection.

The event handling support of DB2CachedRowSet can be used in graphical user interface (GUI) applications to control displays, for logging information about changes to the RowSet as they are made, or to find information about changes to sources other than RowSets. See Example: DB2JdbcRowSet events for details.

For specific details on working with DB2CachedRowSets, see the following topics:

- Create and populate a DB2CachedRowSet
- Access DB2CachedRowSet data and cursor manipulation
- Change DB2CachedRowSet data and reflect changes back to the data source
- Other DB2CachedRowSet features

For information on the event model and event handling, see DB2JdbcRowSet as this support works identically for both types of RowSets. <<

Create and populate a DB2CachedRowSet: >> There are several ways to place data into a DB2CachedRowSet:

- Use the populate method
- Use DB2CachedRowSet properties and DataSources
- Use DB2CachedRowSet properties and JDBC URLs
- Use the setConnection(Connection) method to use an existing database connection
- Use the execute(Connection) method to use an existing database connection
- Use the execute(int) method to group database requests

<H7>Use the populate method

DB2CachedRowSets have a populate method that can be used to put data into the RowSet from a DB2ResultSet object. The following is an example of this approach.

Example: Use the populate method

Note: Read the Code example disclaimer for important legal information.

```
// Establish a connection to the database.
Connection conn = DriverManager.getConnection("jdbc:db2:*local");

// Create a statement and use it to perform a query.
Statement stmt = conn.createStatement();
ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");

// Create and populate a DB2CachedRowSet from it.
DB2CachedRowSet crs = new DB2CachedRowSet();
crs.populate(rs);

// Note: Disconnect the ResultSet, Statement,
// and Connection used to create the RowSet.
rs.close();
stmt.close();
conn.close();

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

crs.close();
```

<H7>Use DB2CachedRowSet properties and DataSources

DB2CachedRowSets have properties that allow the DB2CachedRowSets to accept an SQL query and a DataSource name. They then use the SQL query and DataSource name to create data for themselves. The following is an example of this approach. The reference to the DataSource named BaseDataSource is assumed to be a valid DataSource that has been previously set up.

Example: Use DB2CachedRowSet properties and DataSources

Note: Read the Code example disclaimer for important legal information.

```
// Create a new DB2CachedRowSet
DB2CachedRowSet crs = new DB2CachedRowSet();

// Set the properties that are needed for
// the RowSet to use a DataSource to populate itself.
crs.setDataSourceName("BaseDataSource");
crs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the DataSource and SQL query
// specified to populate itself with data. Once
// the RowSet is populated, it disconnects from the database.
crs.execute();

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

// Eventually, close the RowSet.
crs.close();
```

<H7>Use DB2CachedRowSet properties and JDBC URLs

DB2CachedRowSets have properties that allow the DB2CachedRowSets to accept an SQL query and a JDBC URL. They then use the query and JDBC URL to create data for themselves. The following is an example of this approach.

Example: Use DB2CachedRowSet properties and JDBC URLs

Note: Read the Code example disclaimer for important legal information.

```
// Create a new DB2CachedRowSet
DB2CachedRowSet crs = new DB2CachedRowSet();

// Set the properties that are needed for
// the RowSet to use a JDBC URL to populate itself.
crs.setUrl("jdbc:db2:*local");
crs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the DataSource and SQL query
// specified to populate itself with data. Once
// the RowSet is populated, it disconnects from the database.
crs.execute();

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

// Eventually, close the RowSet.
crs.close();
```

<H7>Use the setConnection(Connection) method to use an existing database connection

To promote the reuse of JDBC Connection objects, the DB2CachedRowSet provides a mechanism for passing an established Connection object to the DB2CachedRowSet that is used to populate the RowSet. If a user-supplied Connection object is passed in, the DB2CachedRowSet does not disconnect it after populating itself.

Example: Use setConnection(Connection) method to use an existing database connection

Note: Read the Code example disclaimer for important legal information.

```
// Establish a JDBC connection to the database.
Connection conn = DriverManager.getConnection("jdbc:db2:*local");

// Create a new DB2CachedRowSet
DB2CachedRowSet crs = new DB2CachedRowSet();

// Set the properties that are needed for the
// RowSet to use an already connected connection
// to populate itself.
crs.setConnection(conn);
crs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the connection that it was provided
// with previously. Once the RowSet is populated, it does not
// close the user-supplied connection.
crs.execute();

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

// Eventually, close the RowSet.
crs.close();
```

<H7>Use the execute(Connection) method to use an existing database connection

To promote the reuse of JDBC Connection objects, the DB2CachedRowSet provides a mechanism for passing an established Connection object to the DB2CachedRowSet when the execute method is called. If a user-supplied Connection object is passed in, the DB2CachedRowSet does not disconnect it after populating itself.

Example: Use execute(Connection) method to use an existing database connection

Note: Read the Code example disclaimer for important legal information.

```
// Establish a JDBC connection to the database.
Connection conn = DriverManager.getConnection("jdbc:db2:*local");

// Create a new DB2CachedRowSet
DB2CachedRowSet crs = new DB2CachedRowSet();

// Set the SQL statement that is to be used to
// populate the RowSet.
crs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method, passing in the connection
// that should be used. Once the Rowset is populated, it does not
// close the user-supplied connection.
crs.execute(conn);

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

// Eventually, close the RowSet.
crs.close();
```

<H7>Use the execute(int) method to group database requests

To reduce the database's workload, the `DB2CachedRowSet` provides a mechanism for grouping SQL statements for several threads into one processing request for the database.

Example: Use `execute(int)` method to group database requests

Note: Read the Code example disclaimer for important legal information.

```
// Create a new DB2CachedRowSet
DB2CachedRowSet crs = new DB2CachedRowSet();

// Set the properties that are needed for
// the RowSet to use a DataSource to populate itself.
crs.setDataSourceName("BaseDataSource");
crs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the DataSource and SQL query
// specified to populate itself with data. Once
// the RowSet is populated, it disconnects from the database.
// This version of the execute method accepts the number of seconds
// that it is willing to wait for its results. By
// allowing a delay, the RowSet can group the requests
// of several users and only process the request against
// the underlying database once.
crs.execute(5);

// Loop through the data in the RowSet.
while (crs.next()) {
    System.out.println("v1 is " + crs.getString(1));
}

// Eventually, close the RowSet.
crs.close();
```

Access `DB2CachedRowSet` data and cursor manipulation: ➤ RowSets depend on `ResultSet` methods. For many operations such as `DB2CachedRowSet` data access and cursor movement, there is no difference at the application level between using a `ResultSet` and using a `RowSet`.

Access `DB2CachedRowSet` data: RowSets and `ResultSets` access data in the same manner. In the following example, the program creates a table and populates it with various data types using JDBC. Once the table is ready, a `DB2CachedRowSet` is created and populated with the information from the table. The example also uses various get methods of the `RowSet` class.

Example: Access `DB2CachedRowSet` data

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.*;
import java.io.*;
import java.math.*;

public class TestProgram
{
    public static void main(String args[])
    {
        // Register the driver.
        try {
```

```

    Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
}
catch (ClassNotFoundException ex) {
    System.out.println("ClassNotFoundException: " +
        ex.getMessage());
    // No need to go any further.
    System.exit(1);
}

try {
    Connection conn = DriverManager.getConnection("jdbc:db2:*local");

    Statement stmt = conn.createStatement();

    // Clean up previous runs
    try {
        stmt.execute("drop table cujosql.test_table");
    }
    catch (SQLException ex) {
        System.out.println("Caught drop table: " + ex.getMessage());
    }

    // Create test table
    stmt.execute("Create table cujosql.test_table (col1 smallint, col2 int, " +
        "col3 bigint, col4 real, col5 float, col6 double, col7 numeric, " +
        "col8 decimal, col9 char(10), col10 varchar(10), col11 date, " +
        "col12 time, col13 timestamp)");
    System.out.println("Table created.");

    // Insert some test rows
    stmt.execute("insert into cujosql.test_table values (1, 1, 1, 1.5, 1.5, 1.5, 1.5, 1.5, 'one', 'one',
        {d '2001-01-01'}, {t '01:01:01'}, {ts '1998-05-26 11:41:12.123456'})");

    stmt.execute("insert into cujosql.test_table values (null, null, null, null, null, null, null, null,
        null, null, null, null, null)");
    System.out.println("Rows inserted");

    ResultSet rs = stmt.executeQuery("select * from cujosql.test_table");
    System.out.println("Query executed");

    // Create a new rowset and populate it...
    DB2CachedRowSet crs = new DB2CachedRowSet();
    crs.populate(rs);
    System.out.println("RowSet populated.");

    conn.close();
    System.out.println("RowSet is detached...");

    System.out.println("Test with getObject");
    int count = 0;
    while (crs.next()) {
        System.out.println("Row " + (++count));
        for (int i = 1; i <= 13; i++) {
            System.out.println(" Col " + i + " value " + crs.getObject(i));
        }
    }

    System.out.println("Test with getXXX... ");
    crs.first();
    System.out.println("Row 1");
    System.out.println(" Col 1 value " + crs.getShort(1));
    System.out.println(" Col 2 value " + crs.getInt(2));
    System.out.println(" Col 3 value " + crs.getLong(3));
    System.out.println(" Col 4 value " + crs.getFloat(4));
    System.out.println(" Col 5 value " + crs.getDouble(5));
    System.out.println(" Col 6 value " + crs.getDouble(6));
    System.out.println(" Col 7 value " + crs.getBigDecimal(7));
}

```

```

        System.out.println(" Col 8 value " + crs.getBigDecimal(8));
        System.out.println(" Col 9 value " + crs.getString(9));
        System.out.println(" Col 10 value " + crs.getString(10));
        System.out.println(" Col 11 value " + crs.getDate(11));
        System.out.println(" Col 12 value " + crs.getTime(12));
        System.out.println(" Col 13 value " + crs.getTimestamp(13));
        crs.next();
        System.out.println("Row 2");
        System.out.println(" Col 1 value " + crs.getShort(1));
        System.out.println(" Col 2 value " + crs.getInt(2));
        System.out.println(" Col 3 value " + crs.getLong(3));
        System.out.println(" Col 4 value " + crs.getFloat(4));
        System.out.println(" Col 5 value " + crs.getDouble(5));
        System.out.println(" Col 6 value " + crs.getDouble(6));
        System.out.println(" Col 7 value " + crs.getBigDecimal(7));
        System.out.println(" Col 8 value " + crs.getBigDecimal(8));
        System.out.println(" Col 9 value " + crs.getString(9));
        System.out.println(" Col 10 value " + crs.getString(10));
        System.out.println(" Col 11 value " + crs.getDate(11));
        System.out.println(" Col 12 value " + crs.getTime(12));
        System.out.println(" Col 13 value " + crs.getTimestamp(13));

        crs.close();
    }
    catch (Exception ex) {
        System.out.println("SQLException: " + ex.getMessage());
        ex.printStackTrace();
    }
}
}

```

Cursor manipulation: RowSets are scrollable and act exactly like a scrollable ResultSet. In the following example, the program creates a table and populates it with data using JDBC. Once the table is ready, a DB2CachedRowSet object is created and is populated with the information from the table. The example also uses various cursor manipulation functions.

Example: Cursor manipulation

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.DB2CachedRowSet;

public class RowSetSample1
{
    public static void main(String args[])
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        }
        catch (ClassNotFoundException ex) {
            System.out.println("ClassNotFoundException: " +
                ex.getMessage());
            // No need to go any further.
            System.exit(1);
        }
    }

    try {
        Connection conn = DriverManager.getConnection("jdbc:db2:*local");

        Statement stmt = conn.createStatement();

        // Clean up previous runs
        try {
            stmt.execute("drop table cujosql.test_table");
        }
    }
}

```

```

catch (SQLException ex) {
    System.out.println("Caught drop table: " + ex.getMessage());
}

// Create a test table
stmt.execute("Create table cujosql.test_table (col1 smallint)");
System.out.println("Table created.");

// Insert some test rows
for (int i = 0; i < 10; i++) {
    stmt.execute("insert into cujosql.test_table values (" + i + ")");
}
System.out.println("Rows inserted");

ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");
System.out.println("Query executed");

// Create a new rowset and populate it...
DB2CachedRowSet crs = new DB2CachedRowSet();
crs.populate(rs);
System.out.println("RowSet populated.");

conn.close();
System.out.println("RowSet is detached...");

System.out.println("Use next()");
while (crs.next()) {
    System.out.println("v1 is " + crs.getShort(1));
}

System.out.println("Use previous()");
while (crs.previous()) {
    System.out.println("value is " + crs.getShort(1));
}

System.out.println("Use relative()");
crs.next();
crs.relative(9);
System.out.println("value is " + crs.getShort(1));

crs.relative(-9);
System.out.println("value is " + crs.getShort(1));

System.out.println("Use absolute()");
crs.absolute(10);
System.out.println("value is " + crs.getShort(1));
crs.absolute(1);
System.out.println("value is " + crs.getShort(1));
crs.absolute(-10);
System.out.println("value is " + crs.getShort(1));
crs.absolute(-1);
System.out.println("value is " + crs.getShort(1));

System.out.println("Test beforeFirst()");
crs.beforeFirst();
System.out.println("isBeforeFirst is " + crs.isBeforeFirst());
crs.next();
System.out.println("move one... isFirst is " + crs.isFirst());

System.out.println("Test afterLast()");
crs.afterLast();
System.out.println("isAfterLast is " + crs.isAfterLast());
crs.previous();
System.out.println("move one... isLast is " + crs.isLast());

System.out.println("Test getRow()");
crs.absolute(7);

```




```

        System.out.println("row should be (7) and is " + crs.getRow() + " value should be (6) and is " + crs.getShort(1));
    }
    crs.close();
}
catch (SQLException ex) {
    System.out.println("SQLException: " + ex.getMessage());
}
}
}
}

```



Change DB2CachedRowSet data and reflect changes back to the data source:  The DB2CachedRowSet uses the same methods as the standard ResultSet interface for making changes to the data in the RowSet object. There is no difference at the application level between changing the data of a RowSet and changing the data of a ResultSet. The DB2CachedRowSet provides the acceptChanges method that is used to reflect changes to the RowSet back to the database where the data came from.

Delete, insert, and update rows in a DB2CachedRowSet: DB2CachedRowSets can be updated. In the following example, the program creates a table and populates it with data using JDBC. Once the table is ready, a DB2CachedRowSet is created and is populated with the information from the table. The example also uses various methods that can be used to update the RowSet and shows how the use of the showDeleted property that allows the application to fetch rows even after they have been deleted. Further, the cancelRowInsert and cancelRowDelete methods are used in the example to allow row insertion or deletion to be undone.

Example: Delete, insert, and update rows in a DB2CachedRowSet

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.DB2CachedRowSet;

public class RowSetSample2
{
    public static void main(String args[])
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        }
        catch (ClassNotFoundException ex) {
            System.out.println("ClassNotFoundException: " +
                ex.getMessage());

            // No need to go any further.
            System.exit(1);
        }

        try {
            Connection conn = DriverManager.getConnection("jdbc:db2:*local");

            Statement stmt = conn.createStatement();

            // Clean up previous runs
            try {
                stmt.execute("drop table cujosql.test_table");
            }

            catch (SQLException ex) {
                System.out.println("Caught drop table: " + ex.getMessage());
            }
        }
    }
}

```

```

// Create test table
stmt.execute("Create table cujosql.test_table (col1 smallint)");
System.out.println("Table created.");

// Insert some test rows
for (int i = 0; i < 10; i++) {
    stmt.execute("insert into cujosql.test_table values (" + i + ")");
}
System.out.println("Rows inserted");

ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");
System.out.println("Query executed");

// Create a new rowset and populate it...
DB2CachedRowSet crs = new DB2CachedRowSet();
crs.populate(rs);
System.out.println("RowSet populated.");

conn.close();
System.out.println("RowSet is detached...");

System.out.println("Delete the first three rows");
crs.next();
crs.deleteRow();
crs.next();
crs.deleteRow();
crs.next();
crs.deleteRow();

crs.beforeFirst();
System.out.println("Insert the value -10 into the RowSet");
crs.moveToInsertRow();
crs.updateShort(1, (short)-10);
crs.insertRow();
crs.moveToCurrentRow();

System.out.println("Update the rows to be the negative of what they now are");
crs.beforeFirst();
while (crs.next())
    short value = crs.getShort(1);
    value = (short)-value;
    crs.updateShort(1, value);
    crs.updateRow();
}

crs.setShowDeleted(true);

System.out.println("RowSet is now (value - inserted - updated - deleted)");
crs.beforeFirst();
while (crs.next()) {
    System.out.println("value is " + crs.getShort(1) + " " +
        crs.rowInserted() + " " +
        crs.rowUpdated() + " " +
        crs.rowDeleted());
}

System.out.println("getShowDeleted is " + crs.getShowDeleted());

System.out.println("Now undo the inserts and deletes");
crs.beforeFirst();
crs.next();
crs.cancelRowDelete();
crs.next();
crs.cancelRowDelete();
crs.next();
crs.cancelRowDelete();

```

```

while (!crs.isLast()) {
    crs.next();
}

crs.cancelRowInsert();

crs.setShowDeleted(false);

System.out.println("RowSet is now (value - inserted - updated - deleted)");
crs.beforeFirst();
while (crs.next()) {
    System.out.println("value is " + crs.getShort(1) + " " +
        crs.rowInserted() + " " +
        crs.rowUpdated() + " " +
        crs.rowDeleted());
}

System.out.println("finally show that calling cancelRowUpdates works");
crs.first();
crs.updateShort(1, (short) 1000);
crs.cancelRowUpdates();
crs.updateRow();
System.out.println("value of row is " + crs.getShort(1));
System.out.println("getShowDeleted is " + crs.getShowDeleted());

crs.close();
}

catch (SQLException ex) {
    System.out.println("SQLException: " + ex.getMessage());
}
}
}

```

Reflect changes to a DB2CachedRowSet back to the underlying database: Once changes have been made to a DB2CachedRowSet, they only exist as long as the RowSet object exists. That is, making changes to a disconnected RowSet has no effect on the database. To reflect the changes of a RowSet in the underlying database, the `acceptChanges` method is used. This method tells the disconnected RowSet to re-establish a connection to the database and attempt to make the changes that have been made to the RowSet to the underlying database. If the changes cannot be safely made to the database due to conflicts with other database changes after the RowSet was created, an exception is thrown and the transaction is rolled back.

Example: Reflect changes to a DB2CachedRowSet back to the underlying database

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.DB2CachedRowSet;

public class RowSetSample3
{
    public static void main(String args[])
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        }
        catch (ClassNotFoundException ex) {
            System.out.println("ClassNotFoundException: " +
                ex.getMessage());
            // No need to go any further.
            System.exit(1);
        }
    }
}

```

```

}

try {
    Connection conn = DriverManager.getConnection("jdbc:db2:*local");

    Statement stmt = conn.createStatement();

    // Clean up previous runs
    try {
        stmt.execute("drop table cujosql.test_table");
    }
    catch (SQLException ex) {
        System.out.println("Caught drop table: " + ex.getMessage());
    }

    // Create test table
    stmt.execute("Create table cujosql.test_table (col1 smallint)");
    System.out.println("Table created.");

    // Insert some test rows
    for (int i = 0; i < 10; i++) {
        stmt.execute("insert into cujosql.test_table values (" + i + ")");
    }
    System.out.println("Rows inserted");

    ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");
    System.out.println("Query executed");

    // Create a new rowset and populate it...
    DB2CachedRowSet crs = new DB2CachedRowSet();
    crs.populate(rs);
    System.out.println("RowSet populated.");

    conn.close();
    System.out.println("RowSet is detached...");

    System.out.println("Delete the first three rows");
    crs.next();
    crs.deleteRow();
    crs.next();
    crs.deleteRow();
    crs.next();
    crs.deleteRow();

    crs.beforeFirst();
    System.out.println("Insert the value -10 into the RowSet");
    crs.moveToInsertRow();
    crs.updateShort(1, (short)-10);
    crs.insertRow();
    crs.moveToCurrentRow();

    System.out.println("Update the rows to be the negative of what they now are");
    crs.beforeFirst();
    while (crs.next()) {
        short value = crs.getShort(1);
        value = (short)-value;
        crs.updateShort(1, value);
        crs.updateRow();
    }

    System.out.println("Now accept the changes to the database");

    crs.setUrl("jdbc:db2:*local");
    crs.setTableName("cujosql.test_table");

    crs.acceptChanges();
    crs.close();
}

```

```

        System.out.println("And the database table looks like this:");
        conn = DriverManager.getConnection("jdbc:db2:localhost");
        stmt = conn.createStatement();
        rs = stmt.executeQuery("select col1 from cujosql.test_table");
        while (rs.next()) {
            System.out.println("Value from table is " + rs.getShort(1));
        }

        conn.close();
    }
    catch (SQLException ex) {
        System.out.println("SQLException: " + ex.getMessage());
    }
}
}
}

```



Other DB2CachedRowSet features: ➤ In addition to working like a ResultSet as several examples have shown, the DB2CachedRowSet class has some additional functionality that makes it more flexible to use. Methods are provided for turning either the entire Java^(TM) Database Connectivity (JDBC) RowSet or just a portion of it into a Java collection. Moreover, because of their disconnected nature, DB2CachedRowSets do not have a strict one-to-one relationship with ResultSets.

With the methods provided by DB2CachedRowSet, you can perform the following tasks:

- Obtain collections from DB2CachedRowSets
- Create copies of RowSets
- Create shares for RowSets

Obtain collections from DB2CachedRowSets: There are three methods that return some form of a collection from a DB2CachedRowSet object. They are the following:

- **toCollection** returns an ArrayList (that is, one entry for each row) of vectors (that is, one entry for each column).
- **toCollection(int columnIndex)** returns a vector containing the value for each row from the given column.
- **getColumn(int columnIndex)** returns an array containing the value for each column for a given column.

The major difference between toCollection(int columnIndex) and getColumn(int columnIndex) is that the getColumn method can return an array of primitive types. Therefore, if columnIndex represents a column that has integer data, an integer array is returned and not an array containing java.lang.Integer objects.

The following example shows how you can use these methods.

Example: Obtain collections from DB2CachedRowSets

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.DB2CachedRowSet;
import java.util.*;

public class RowSetSample4
{
    public static void main(String args[])
    {

```

```

// Register the driver.
try {
    Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
}
catch (ClassNotFoundException ex) {
    System.out.println("ClassNotFoundException: " +
        ex.getMessage());
    // No need to go any further.
    System.exit(1);
}

try {
    Connection conn = DriverManager.getConnection("jdbc:db2:*local");
    Statement stmt = conn.createStatement();

    // Clean up previous runs
    try {
        stmt.execute("drop table cujosql.test_table");
    }
    catch (SQLException ex) {
        System.out.println("Caught drop table: " + ex.getMessage());
    }

    // Create test table
    stmt.execute("Create table cujosql.test_table (col1 smallint, col2 smallint)");
    System.out.println("Table created.");

    // Insert some test rows
    for (int i = 0; i < 10; i++) {
        stmt.execute("insert into cujosql.test_table values (" + i + ", " + (i + 100) + ")");
    }
    System.out.println("Rows inserted");

    ResultSet rs = stmt.executeQuery("select * from cujosql.test_table");
    System.out.println("Query executed");

    // Create a new rowset and populate it...
    DB2CachedRowSet crs = new DB2CachedRowSet();
    crs.populate(rs);
    System.out.println("RowSet populated.");

    conn.close();
    System.out.println("RowSet is detached...");

    System.out.println("Test the toCollection() method");
    Collection collection = crs.toCollection();
    ArrayList map = (ArrayList) collection;

    System.out.println("size is " + map.size());
    Iterator iter = map.iterator();
    int row = 1;
    while (iter.hasNext()) {
        System.out.print("row [" + (row++) + "]: \t");

        Vector vector = (Vector)iter.next();
        Iterator innerIter = vector.iterator();
        int i = 1;
        while (innerIter.hasNext()) {
            System.out.print(" [" + (i++) + "]= " + innerIter.next() + "; \t");
        }
        System.out.println();
    }
    System.out.println("Test the toCollection(int) method");
    collection = crs.toCollection(2);
    Vector vector = (Vector) collection;

    iter = vector.iterator();

```

```

while (iter.hasNext()) {
    System.out.println("Iter: Value is " + iter.next());
}

System.out.println("Test the getColumn(int) method");
Object values = crs.getColumn(2);
short[] shorts = (short [])values;

for (int i =0; i < shorts.length; i++) {
    System.out.println("Array: Value is " + shorts[i]);
}
}
catch (SQLException ex) {
    System.out.println("SQLException: " + ex.getMessage());
}
}
}

```

Create copies of RowSets: The createCopy method creates a copy of the DB2CachedRowSet. All the data associated with the RowSet is replicated along with all control structures, properties, and status flags.

The following example shows how you can use this method.

Example: Create copies of RowSets

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.*;
import java.io.*;

public class RowSetSample5
{
    public static void main(String args[])
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        }
        catch (ClassNotFoundException ex) {
            System.out.println("ClassNotFoundException: " +
                ex.getMessage());
            // No need to go any further.
            System.exit(1);
        }

        try {
            Connection conn = DriverManager.getConnection("jdbc:db2:*local");

            Statement stmt = conn.createStatement();

            // Clean up previous runs
            try {
                stmt.execute("drop table cujosql.test_table");
            }
            catch (SQLException ex) {
                System.out.println("Caught drop table: " + ex.getMessage());
            }

            // Create test table
            stmt.execute("Create table cujosql.test_table (col1 smallint)");
            System.out.println("Table created.");

            // Insert some test rows

```

```

for (int i = 0; i < 10; i++) {
    stmt.execute("insert into cujosql.test_table values (" + i + ")");
}
System.out.println("Rows inserted");

ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");
System.out.println("Query executed");

// Create a new rowset and populate it...
DB2CachedRowSet crs = new DB2CachedRowSet();
crs.populate(rs);
System.out.println("RowSet populated.");

conn.close();
System.out.println("RowSet is detached...");

System.out.println("Now some new RowSets from one.");
DB2CachedRowSet crs2 = crs.createCopy();
DB2CachedRowSet crs3 = crs.createCopy();

System.out.println("Change the second one to be negated values");
crs2.beforeFirst();
while (crs2.next()) {
    short value = crs2.getShort(1);
    value = (short)-value;
    crs2.updateShort(1, value);
    crs2.updateRow();
}

crs.beforeFirst();
crs2.beforeFirst();
crs3.beforeFirst();
System.out.println("Now look at all three of them again");

while (crs.next()) {
    crs2.next();
    crs3.next();
    System.out.println("Values: crs: " + crs.getShort(1) + ", crs2: " + crs2.getShort(1) +
        ", crs3: " + crs3.getShort(1));
}
}
catch (Exception ex) {
    System.out.println("SQLException: " + ex.getMessage());
    ex.printStackTrace();
}
}
}

```

Create shares for RowSets: The createShared method creates a new RowSet object with high-level status information and allows two RowSet objects to share the same underlying physical data.

The following example shows how you can use this method.

Example: Create shares of RowSets

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.*;
import java.io.*;

public class RowSetSample5
{

```



```

public static void main(String args[])
{
    // Register the driver.
    try {
        Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
    }
    catch (ClassNotFoundException ex) {
        System.out.println("ClassNotFoundException: " +
            ex.getMessage());
        // No need to go any further.
        System.exit(1);
    }

    try {
        Connection conn = DriverManager.getConnection("jdbc:db2:*local");

        Statement stmt = conn.createStatement();

        // Clean up previous runs
        try {
            stmt.execute("drop table cujosql.test_table");
        }
        catch (SQLException ex) {
            System.out.println("Caught drop table: " + ex.getMessage());
        }

        // Create test table
        stmt.execute("Create table cujosql.test_table (col1 smallint)");
        System.out.println("Table created.");

        // Insert some test rows
        for (int i = 0; i < 10; i++) {
            stmt.execute("insert into cujosql.test_table values (" + i + ")");
        }
        System.out.println("Rows inserted");

        ResultSet rs = stmt.executeQuery("select col1 from cujosql.test_table");
        System.out.println("Query executed");

        // Create a new rowset and populate it...
        DB2CachedRowSet crs = new DB2CachedRowSet();
        crs.populate(rs);
        System.out.println("RowSet populated.");

        conn.close();
        System.out.println("RowSet is detached...");

        System.out.println("Test the createShared functionality (create 2 shares)");
        DB2CachedRowSet crs2 = crs.createShared();
        DB2CachedRowSet crs3 = crs.createShared();

        System.out.println("Use the original to update value 5 of the table");
        crs.absolute(5);
        crs.updateShort(1, (short)-5);
        crs.updateRow();

        crs.beforeFirst();
        crs2.afterLast();

        System.out.println("Now move the cursors in opposite directions of the same data.");

        while (crs.next()) {
            crs2.previous();
            crs3.next();
            System.out.println("Values: crs: " + crs.getShort(1) + ", crs2: " + crs2.getShort(1) +
                ", crs3: " + crs3.getShort(1));
        }
    }
}

```

```

        crs.close();
        crs2.close();
        crs3.close();
    }
    catch (Exception ex) {
        System.out.println("SQLException: " + ex.getMessage());
        ex.printStackTrace();
    }
}
}
}

```



DB2JdbcRowSet: ➤ The DB2JdbcRowSet is a connected RowSet, meaning that it can only be used with the support of an underlying Connection object, PreparedStatement object, or ResultSet object. Its implementation adheres closely to the description of a JdbcRowSet.

Use DB2JdbcRowSet: Because the DB2JdbcRowSet object supports events described in the Java^(TM) Database Connectivity (JDBC) 3.0 specification for all RowSets, it can serve as an intermediate object between a local database and other objects that must be notified about changes to the database data.

As an example, assume that you are working in an environment where you have a main database and several Personal Digital Assistants (PDAs) that use a wireless protocol to connect to it. A DB2JdbcRowSet object can be used to move to a row and update it by using a master application that is running on the server. The row update causes an event to be generated by the RowSet component. If there is a service running that is responsible for sending out updates to the PDAs, it can register itself as a "listener" of the RowSet. Each time that it receives a RowSet event, it can generate the appropriate update and send it out to the wireless devices.

Refer to Example: DB2JdbcRowSet events for more information.

Create JDBCRowSets: There are several methods provided for creating a DB2JDBCRowSet object. Each is outlined as follows.

<h7>Use DB2JdbcRowSet properties and DataSources

DB2JdbcRowSets have properties that accept an SQL query and a DataSource name. The DB2JdbcRowSets are then ready to be used. The following is an example of this approach. The reference to the DataSource named BaseDataSource is assumed to be a valid DataSource that has been previously set up.

Example: Use DB2JdbcRowSet properties and DataSources

Note: Read the Code example disclaimer for important legal information.

```

// Create a new DB2JdbcRowSet
DB2JdbcRowSet jrs = new DB2JdbcRowSet();

// Set the properties that are needed for
// the RowSet to be processed.
jrs.setDataSourceName("BaseDataSource");
jrs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This method causes
// the RowSet to use the DataSource and SQL query
// specified to prepare itself for data processing.
jrs.execute();

// Loop through the data in the RowSet.
while (jrs.next()) {
    System.out.println("v1 is " + jrs.getString(1));
}

```

```

}

// Eventually, close the RowSet.
jrs.close();

```

Use DB2JdbcRowSet properties and JDBC URLs

DB2JdbcRowSets have properties that accept an SQL query and a JDBC URL. The DB2JdbcRowSets are then ready to be used. The following is an example of this approach:

Example: Use DB2JdbcRowSet properties and JDBC URLs

Note: Read the Code example disclaimer for important legal information.

```

// Create a new DB2JdbcRowSet
DB2JdbcRowSet jrs = new DB2JdbcRowSet();

// Set the properties that are needed for
// the RowSet to be processed.
jrs.setUrl("jdbc:db2:*local");
jrs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the URL and SQL query specified
// previously to prepare itself for data processing.
jrs.execute();

// Loop through the data in the RowSet.
while (jrs.next()) {
    System.out.println("v1 is " + jrs.getString(1));
}

// Eventually, close the RowSet.
jrs.close();

```

Use the setConnection(Connection) method to use an existing database connection

To promote the reuse of JDBC Connection objects, the DB2JdbcRowSet allows you to pass an established connection to the DB2JdbcRowSet. This connection is used by the DB2JdbcRowSet to prepare itself for usage when the execute method is called.

Example: Use the setConnection method

Note: Read the Code example disclaimer for important legal information.

```

// Establish a JDBC Connection to the database.
Connection conn = DriverManager.getConnection("jdbc:db2:*local");

// Create a new DB2JdbcRowSet.
DB2JdbcRowSet jrs = new DB2JdbcRowSet();

// Set the properties that are needed for
// the RowSet to use an established connection.
jrs.setConnection(conn);
jrs.setCommand("select col1 from cujosql.test_table");

// Call the RowSet execute method. This causes
// the RowSet to use the connection that it was provided
// previously to prepare itself for data processing.
jrs.execute();

// Loop through the data in the RowSet.
while (jrs.next()) {
    System.out.println("v1 is " + jrs.getString(1));
}

```

```

}

// Eventually, close the RowSet.
jrs.close();

```

Access data and cursor movement: Manipulation of the cursor position and access to the database data through a DB2JdbcRowSet are handled by the underlying ResultSet object. Tasks that can be done with a ResultSet object also apply to the DB2JdbcRowSet object.

Change data and reflecting changes to the underlying database: Support for updating the database through a DB2JdbcRowSet is handled completely by the underlying ResultSet object. Tasks that can be done with a ResultSet object also apply to the DB2JdbcRowSet object. <<

DB2JdbcRowSet events: >> All RowSet implementations support event handling for situations that are of interest to other components. This support allows application components to "talk" to each other when events happen to them. For example, updating a database row through a RowSet can cause a Graphical User Interface (GUI) table shown to you to update itself.

In the following example, the main method does the update to the RowSet and is your core application. The listener is part of your wireless server used by your disconnected clients in the field. It is possible to tie these two aspects of a business together without getting the code for the two processes intermingled. While the event support of RowSets was designed primarily for updating GUIs with database data, it works perfectly for this type of application problem.

Example: DB2JdbcRowSet events

Note: Read the Code example disclaimer for important legal information.

```

import java.sql.*;
import javax.sql.*;
import com.ibm.db2.jdbc.app.DB2JdbcRowSet;

public class RowSetEvents {
    public static void main(String args[])
    {
        // Register the driver.
        try {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver");
        } catch (ClassNotFoundException ex) {
            System.out.println("ClassNotFoundException: " +
                ex.getMessage());
            // No need to go any further.
            System.exit(1);
        }

        try {
            // Obtain the JDBC Connection and Statement needed to set
            // up this example.
            Connection conn = DriverManager.getConnection("jdbc:db2:*local");
            Statement stmt = conn.createStatement();

            // Clean up any previous runs.
            try {
                stmt.execute("drop table cujosql.test_table");
            } catch (SQLException ex) {
                System.out.println("Caught drop table: " + ex.getMessage());
            }

            // Create the test table
            stmt.execute("Create table cujosql.test_table (col1 smallint)");
            System.out.println("Table created.");
        }
    }
}

```

```

// Populate the table with data.
for (int i = 0; i < 10; i++) {
    stmt.execute("insert into cujosql.test_table values (" + i + ")");
}
System.out.println("Rows inserted");

// Remove the setup objects.
stmt.close();
conn.close();

// Create a new rowset and set the properties need to
// process it.
DB2JdbcRowSet jrs = new DB2JdbcRowSet();
jrs.setUrl("jdbc:db2:*local");
jrs.setCommand("select col1 from cujosql.test_table");
jrs.setConcurrency(ResultSet.CONCUR_UPDATEABLE);

// Give the RowSet object a listener. This object handles
// special processing when certain actions are done on
// the RowSet.
jrs.addRowSetListener(new MyListener());

// Process the RowSet to provide access to the database data.
jrs.execute();

// Cause a few cursor change events. These events cause the cursorMoved
// method in the listener object to get control.
jrs.next();
jrs.next();
jrs.next();

// Cause a row change event to occur. This event causes the rowChanged method
// in the listener object to get control.
jrs.updateShort(1, (short)6);
jrs.updateRow();

// Finally, cause a RowSet change event to occur. This causes the
// rowSetChanged method in the listener object to get control.
jrs.execute();

// When completed, close the RowSet.
jrs.close();
} catch (SQLException ex) {
    ex.printStackTrace();
}
}

/**
 * This is an example of a listener. This example prints messages that show
 * how control flow moves through the application and offers some
 * suggestions about what might be done if the application were fully implemented.
 */
class MyListener
implements RowSetListener {
    public void cursorMoved(RowSetEvent rse) {
        System.out.println("Event to do: Cursor position changed.");
        System.out.println(" For the remote system, do nothing ");
        System.out.println(" when this event happened. The remote view of the data");
        System.out.println(" could be controlled separately from the local view.");
        try {
            DB2JdbcRowSet rs = (DB2JdbcRowSet) rse.getSource();
            System.out.println("row is " + rs.getRow() + ". \n\n");
        } catch (SQLException e) {
            System.out.println("To do: Properly handle possible problems.");
        }
    }
}

```

```

}

public void rowChanged(RowSetEvent rse) {
    System.out.println("Event to do: Row changed.");
    System.out.println(" Tell the remote system that a row has changed. Then,");
    System.out.println(" pass all the values only for that row to the ");
    System.out.println(" remote system.");
    try {
        DB2JdbcRowSet rs = (DB2JdbcRowSet) rse.getSource();
        System.out.println("new values are " + rs.getShort(1) + ". \n\n");
    } catch (SQLException e) {
        System.out.println("To do: Properly handle possible problems.");
    }
}

public void rowSetChanged(RowSetEvent rse) {
    System.out.println("Event to do: RowSet changed.");
    System.out.println(" If there is a remote RowSet already established, ");
    System.out.println(" tell the remote system that the values it ");
    System.out.println(" has should be thrown out. Then, pass all ");
    System.out.println(" the current values to it.\n\n");
}
}

```



Performance tips for the IBM Developer Kit for Java JDBC driver

The IBM Developer Kit for Java^(TM) JDBC driver is designed to be a high performance Java interface for working with the database. However, getting the best possible performance requires that you build your applications in a way that takes advantage of the strengths the JDBC driver has to offer. The following tips are considered good JDBC programming practice. Most are not specific to the native JDBC driver. Therefore, applications written according to these guidelines also perform well if used with JDBC drivers other than the native JDBC driver.

- Avoid SELECT * SQL queries
- Use getXXX(int) instead of getXXX(String)
- Avoid getObject calls for Java primitive types
- Use PreparedStatement over Statement
- Avoid expensive DatabaseMetaData calls
- Use the correct commit level for your application
- Consider storing data in Unicode
- Use stored procedures
- Use BigInt instead of Numeric/Decimal
- Close your JDBC resources explicitly when done with them
- Use connection pooling
- Consider using PreparedStatement pooling
- Use efficient SQL

Avoid SELECT * SQL queries

SELECT * FROM... is a common way to state a query in SQL. Often, however, you do not need to query all the fields. For each column that is to be returned, the JDBC driver must do the additional work of binding and returning the row. Even if your application never uses a particular column, the JDBC driver has to be made aware of it and has to reserve space for its use. If your tables have few columns that are not used, this is not significant overhead. For a large number of unused columns, however, the overhead can be significant. A better solution is to list the columns that your application is interested in individually, like this:

```
SELECT COL1, COL2, COL3 FROM...
```

Use getXXX(int) instead of getXXX(String)

Use the ResultSet getXXX methods that take numeric values instead of the versions that take column names. While the freedom to use your column names instead of numeric constants seems like an advantage, the database itself only knows how to deal with column indexes. Therefore, each getXXX method with a column name you call must be resolved by the JDBC driver before it can be passed to the database. Because getXXX methods are typically called inside loops that could be run millions of times, this little bit of overhead can rapidly accumulate.

Avoid getObject calls for Java primitive types

When getting values from the database of primitive types (ints, longs, floats, and so on), it is faster to use the get method specific to the primitive type (getInt, getLong, getFloat) than to use getObject. The getObject call does the work of the get for the primitive type, and then creates an object to return to you. This is typically done in loops, potentially creating millions of objects with short lifespans. Using getObject for primitive commands has the added drawback of frequently activating the garbage collector, further degrading performance.

Use PreparedStatement over Statement

If you are writing an SQL statement that is used more than once, it performs better as a PreparedStatement than as a Statement object. Every time you run a statement, you go through a two step process: the statement is prepared, and then the statement is processed. When you use a prepared statement, the statement is prepared only at the time that it is constructed, not each time it is run. Though it is recognized that a PreparedStatement performs faster than a Statement, this advantage is often neglected by programmers. Due to the performance boost that PreparedStatements provide, it is wise to use them in the design of your applications wherever possible (see PreparedStatement pooling).

Avoid DatabaseMetaData calls

Be aware that some of the DatabaseMetaData calls can be expensive. In particular, the getBestRowIdentifier, getCrossReference, getExportedKeys, and getImportedKeys methods can be costly. Some DatabaseMetaData calls involve complex join conditions over system-level tables. Use them only if you need their information, not just for convenience.

Use the correct commit level for your application

JDBC provides several commit levels which determine how multiple transactions against the system affect each other (see Transactions for more details). The default is to use the lowest commit level. This means that transactions can see some of each other's work through commit boundaries. This introduces the possibility of certain database anomalies. Some programmers increase the commit level so that they do not have to worry about these anomalies occurring. Be aware that higher commit levels involve the database hanging onto more course-grained locks. This limits the amount of concurrency that the system can have, severely slowing the performance of some applications. Often, the anomaly conditions cannot occur because of the design of the application in the first place. Take time to understand what you are trying to accomplish and limit your transaction isolation level to the lowest level you can safely use.

Consider storing data in Unicode

Java requires all character data that it works with (Strings) to be in Unicode. Therefore, any table that does not have Unicode data requires the JDBC driver to translate the data back and forth as it is put into the database and retrieved out of the database. If the table is already in Unicode, the JDBC driver does not need to translate the data and can therefore place the data from the database faster. Take care to understand that data in Unicode may not work with non-Java applications, which do not know how to deal

with Unicode. Also keep in mind that non-character data does not perform any faster, as there is never a translation of this data. Another consideration is that data stored in Unicode takes up twice as much space as single byte data does. If you have many character columns that are read many times, however, the performance gained by storing your data in Unicode can be significant.

Use stored procedures

The use of stored procedures is supported in Java. Stored procedures can perform faster by allowing the JDBC driver to run static SQL instead of dynamic SQL. Do not create stored procedures for each individual SQL statement you run in your program. Where possible, however, create a stored procedure that runs a group of SQL statements.

Use BigInteger instead of Numeric or Decimal

Instead of using Numeric or Decimal fields that have a scale of 0, use the BigInteger data type. BigInteger translates directly into the Java primitive type Long whereas Numeric or Decimal data types translate into String or BigDecimal objects. As noted in Avoid getObject calls for Java primitive types, using primitive data types is preferable to using types that require object creation.

Explicitly close your JDBC resources when done with them

ResultSets, Statements, and Connections should be explicitly closed by the application when they are no longer needed. This allows the resources to be cleaned up in the most efficient way possible and can increase performance. Further, database resources that are not explicitly closed can cause resource leaks and database locks to be held longer than necessary. This can lead to application failures or reduced concurrency in applications.

Use connection pooling

Connection pooling is a strategy by which JDBC Connection objects get reused for multiple users instead of each user request creating its own Connection object. Connection objects are expensive to create. Instead of having each user create a new one, a pool of them should be shared in applications where performance is critical. Many products (such as WebSphere) provide Connection pooling support that can be used with little additional effort on the user's part. If you do not want to use a product with connection pooling support, or prefer to build your own for better control over how the pool works and performs, it is reasonably easy to do so.

Consider using PreparedStatement pooling


Statement pooling works similarly to Connection pooling. Instead of just putting Connections into a pool, put an object that contains the Connection and the PreparedStatements a pool. Then, retrieve that object and access the specific statement you want to use. This can dramatically increase performance.

Use efficient SQL

Because JDBC is built on top of SQL, just about anything that makes for efficient SQL also makes for efficient JDBC. Hence, JDBC benefits from optimized queries, wisely chosen indices, and other aspects of good SQL design.

Access databases using IBM Developer Kit for Java DB2 SQLJ support

DB2 Structured Query Language for Java^(TM) (SQLJ) support is based on the SQLJ ANSI standard. The DB2 SQLJ support is contained in the IBM Developer Kit for Java. DB2 SQLJ support allows you to create, build, and run embedded SQL for Java applications.

The SQLJ support provided by the IBM Developer Kit for Java includes the SQLJ run-time classes, and is available in /QIBM/ProdData/Java400/ext/runtime.zip. For more information on the SQLJ run-time classes, refer to the Runtime API documentation supplied in the Implementation from www.sqlj.org .

SQLJ tools

The following tools are also included in the SQLJ support provided by the IBM Developer Kit for Java:

- The SQLJ translator, `sqlj`, replaces embedded SQL statements in the SQLJ program with Java source statements and generates a serialized profile that contains information about the SQLJ operations that are found in the SQLJ program.
- The DB2 SQLJ Profile Customizer, `db2profrc`, precompiles the SQL statements stored in the generated profile and generates a package in the DB2 database.
- The DB2 SQLJ Profile Printer, `db2profp`, prints the contents of a DB2 customized profile in plain text.
- The SQLJ profile auditor installer, `profdb`, installs and uninstalls debugging class-auditors into an existing set of binary profiles.
- The SQLJ profile conversion tool, `profconv`, converts a serialized profile instance to Java class format.

Note: These tools must be run in the Qshell Interpreter.

DB2 SQLJ restrictions

When you create DB2 applications with SQLJ, you should be aware of the following restrictions:

- DB2 SQLJ support adheres to standard DB2 Universal Database restrictions on issuing SQL statements.
- The DB2 SQLJ profile customizer should only be run on profiles associated with connections to the local database.
- The SQLJ Reference Implementation requires JDK 1.1, or higher. See Support for multiple Java Development Kits (JDKs) for more information on running multiple versions of the Java Development Kit.

For information on using SQL in your Java applications, see Embed SQL Statements in your Java application and Compile and run SQLJ programs.

Structured Query Language for Java profiles


Profiles are generated by the SQLJ Translator, `sqlj`, when you translate the SQLJ source file. Profiles are serialized binary files. That is why these files have a `.ser` extension. These files contain the SQL statements from the associated SQLJ source file.

To generate profiles from your SQLJ source code, run the SQLJ translator, `sqlj`, on your `.sqlj` file.

For more information, see Compile and run SQLJ programs.

The structured query language for Java (SQLJ) translator (`sqlj`)

The SQLJ translator, `sqlj`, generates a serialized profile containing information about the SQL operations found in the SQLJ program. The SQLJ translator uses the /QIBM/ProdData/Java400/ext/translator.zip file.

For more information on `sqlj` command line options, refer to the SQLJ User's Guide and Reference supplied in the Implementation from www.sqlj.org .

Precompile SQL statements in a profile using the DB2 SQLJ profile customizer, `db2profrc`

You can use the DB2 SQLJ Profile Customizer, `db2profrc`, to make your Java^(TM) application work more efficiently with your database.

The DB2 SQLJ Profile Customizer does the following:

- Precompiles the SQL statements that are stored in a profile and generates a package in the DB2 database.
- Customizes the SQLJ profile by replacing the SQL statements with references to the associated statement in the package that was created.

To precompile the SQL statements in a profile, type in the following at the Qshell command prompt:

```
db2profc MyClass_SJProfile0.ser
```

Where *MyClass_SJProfile0.ser* is the name of the profile you want to precompile.

DB2 SQLJ Profile Customizer usage and syntax

```
db2profc[options] <SQLJ_profile_name>
```

Where *SQLJ_profile_name* is the name of the profile to be printed and *options* is the list of options you want.

The options available for db2profc are the following:

- -URL=<JDBC_URL>
- -user=<username>
- -password=<password>
- -package=<library_name/package_name>
- -commitctrl=<commitment_control>
- -datefmt=<date_format>
- -datesep=<date_separator>
- -timefmt=<time_format>
- -timesep=<time_separator>
- -decimalpt=<decimal_point>
- -stmtCCSID=<CCSID>
- -sorttbl=<library_name/sort_sequence_table_name>
- -langID=<language_identifier>

The following are the descriptions of these options:

-URL=<JDBC_URL>

Where *JDBC_URL* is the URL of the JDBC connection. The syntax for the URL is:

```
"jdbc:db2:systemName"
```

For more information, see Access your iSeries database with the IBM Developer Kit for Java JDBC driver.

-user=<username>

Where *username* is your username. The default value is the user ID of the current user that is signed on for local connection.

-password=<password>

Where *password* is your password. The default value is the password of the current user that is signed on for local connection.

-package=<library name/package name>

Where *library name* is the library where the package is placed, and *package name* is the name of the package to be generated. The default library name is QUSRSYS. The default package name is generated from the name of the profile. The maximum length for the package name is 10 characters. Because the SQLJ profile name is always longer than 10 characters, the default

package name that is constructed is different from the profile name. The default package name is constructed by concatenating the first letters of the profile name with the profile key number. If the profile key number is greater than 10 characters long, then the last 10 characters of the profile key number is used for the default package name. For example, the following chart shows some profile names and their default package names:

Profile name	Default package name
App_SJProfile0	App_SJPro0
App_SJProfile01234	App_S01234
App_SJProfile012345678	A012345678
App_SJProfile01234567891	1234567891

-commitctl=<commitment_control>

Where *commitment_control* is the level of commitment control you want. Commitment control can have any one of the following character values:

Value	Definition
C	*CHG. Dirty reads, nonrepeatable reads and phantom reads are possible.
S	*CS. Dirty reads are not possible, but non-repeatable reads and phantom reads are possible.
A	*ALL. Dirty reads and nonrepeatable reads are not possible, but phantom reads are possible.
N	*NONE. Dirty reads, nonrepeatable reads, and phantom reads are not possible. This is the default.

-datefmt=<date_format>

Where *date_format* is the type of date formatting you want. Date format can have any one of the following values:

Value	Definition
USA	IBM USA standard (mm.dd.yyyy, hh:mm a.m., hh:mm p.m.)
ISO	International Standards Organization (yyyy-mm-dd, hh.mm.ss) This is the default.
EUR	IBM European Standard (dd.mm.yyyy, hh.mm.ss)
JIS	Japanese Industrial Standard Christian Era (yyyy-mm-dd, hh:mm:ss)
MDY	Month/Day/Year (mm/d/yy)
DMY	Day/Month/Year (dd/mm/yy)
YMD	Year/Month/Day (yy/mm/dd)
JUL	Julian (yy/ddd)

Date format is used when accessing date result columns. All output date fields are returned in the specified format. For input date strings, the specified value is used to determine whether the date is specified in a valid format. The default value is ISO.

-datesep=<date_separator>

Where *date_separator* is the type of separator you want to use. Date separator is used when accessing date result columns. Date separator can be any of the following values:

Value	Definition
/	A slash is used.
.	A period is used.
,	A comma is used.
-	A dash is used. This is the default.
blank	A space is used.

-timefmt=<time_format>

Where *time_format* is the format you want to use to display time fields. Time format is used when accessing time result columns. For input time strings, the specified value is used to determine whether the time is specified in a valid format. Time format can be any one of the following values:

Value	Definition
USA	IBM USA standard (mm.dd.yyyy, hh:mm a.m., hh:mm p.m.)
ISO	International Standards Organization (yyyy-mm-dd, hh.mm.ss) This is the default.
EUR	IBM European Standard (dd.mm.yyyy, hh.mm.ss)
JIS	Japanese Industrial Standard Christian Era (yyyy-mm-dd, hh:mm:ss)
HMS	Hour/Minute/Second (hh:mm:ss)

-timesep=<time_separator>

Where *time_separator* is the character you want to use to access your time result columns. Time separator can be any one of the following values:

Value	Definition
:	A colon is used.
.	A period is used. This is the default.
,	A comma is used.
blank	A space is used.

-decimalpt=<decimal_point>

Where *decimal_point* is the decimal point you want to use. The decimal point is used for numeric constants in SQL statements. Decimal point can be any one of the following values:

Value	Definition
.	A period is used. This is the default.
,	A comma is used.

-stmtCCSID=<CCSID>

Where *CCSID* is the coded character set identifier for the SQL statements that are prepared into the package. The value of the job during customization time is the default value.


-sorttbl=<library_name/sort_sequence_table_name>

Where *library_name/sort_sequence_table_name* is the location and table name of the sort sequence table you want to use. The sort sequence table is used for string comparisons in SQL statements. The library name and sort sequence table name each have limits of 10 characters. The default value is taken from the job during customization time.

-langID=<language_identifier>

Where *language identifier* is the language identifier you want to use. The default value for the language identifier is taken from the current job during customization time. The language identifier is used in conjunction with the sort sequence table.

For a more detailed information on any of these fields, see DB2 for iSeries SQL Programming Concepts,

SC41-5611  .

Print the contents of DB2 SQLJ profiles (db2profp and profp)

The DB2 SQLJ Profile Printer, db2profp, prints the contents of a DB2 customized profile in plain text. The Profile Printer, profp, prints the contents of profiles generated by the SQLJ translator in plain text.

To print the content of the profiles generated by the SQLJ translator in plain text, use the profp utility as follows:

```
profp MyClass_SJProfile0.ser
```

Where *MyClass_SJProfile0.ser* is the name of the profile you want to print.

To print the content of the DB2 customized version of the profile in plain text, use the db2profp utility as follows:

```
db2profp MyClass_SJProfile0.ser
```

Where *MyClass_SJProfile0.ser* is the name of the profile you want to print.

Note: If you run db2profp on an uncustomized profile, it tells you that the profile has not been customized. If you run profp on a customized profile, it displays the contents of the profile without the customizations.

DB2 SQLJ Profile Printer usage and syntax:

```
db2profp [options] <SQLJ_profile_name>
```

Where *SQLJ_profile_name* is the name of the profile to be printed and *options* is the list of options you want.

The options available for db2profp are the following:

-URL=<JDBC_URL>

Where *JDBC_URL* is the URL you want to connect to. For more information, see Access your iSeries database with the IBM Developer Kit for Java JDBC driver.

-user=<username>

Where *username* is the user name is your user profile.

-password=<password>

Where *password* is the password of your user profile.

SQLJ profile auditor installer (profdb)

The SQLJ profile auditor installer (profdb) installs and uninstalls debugging class-auditors. The debugging class-auditors are installed into an existing set of binary profiles. Once the debugging class-auditors are installed, all RTStatement and RTResultSet calls made during application run time are logged. They can be logged to a file or standard output. The logs can then be inspected to verify the behavior and trace errors of the application. Note that only the calls made to the underlying RTStatement and RTResultSetcall interface at run time are audited.

To install debugging class-auditors, enter the following at the Qshell command prompt:


```
profdb MyClass_SJProfile0.ser
```

Where *MyClass_SJProfile0.ser* is the name of the profile that was generated by the SQLJ Translator.

To uninstall debugging class-auditors, enter the following at the Qshell command prompt:

```
profdb -Cuninstall MyClass_SJProfile.ser
```

Where *MyClass_SJProfile0.ser* is the name of the profile that was generated by the SQLJ Translator.

For more information on profdb command line options, see www.sqlj.org , select the Implementation category and go to the `sqlj.runtime.profile.util.AuditorInstaller` class that is in the Runtime API documentation.

Convert a serialized profile instance to Java class format using the SQLJ profile conversion tool (profconv)

The SQLJ profile conversion tool (profconv) converts a serialized profile instance to Java^(TM) class format. The profconv tool is needed because some browsers do not support loading a serialized object from a resource file that is associated with an applet. Run the profconv utility to perform the conversion.

To run the profconv utility, type the following on the Qshell command line:

```
profconv MyApp_SJProfile0.ser
```

where *MyApp_SJProfile0.ser* is the name of profile instance you want to convert.

The profconv tool invokes `sqlj -ser2class`. See `sqlj` for command line options.

Embed SQL statements in your Java application

Static SQL statements in SQLJ are in SQLJ clauses. SQLJ clauses begin with `#sql` and end with a semicolon (;) character.

Before you create any SQLJ clauses in your Java^(TM) application, import the following packages:

- `import java.sql.*;`
- `import sqlj.runtime.*;`
- `import sqlj.runtime.ref.*;`

The simplest SQLJ clauses are clauses that can be processed and consist of the token `#sql` followed by an SQL statement enclosed in braces. For example, the following SQLJ clause may appear wherever a Java statement may legally appear:

```
#sql { DELETE FROM TAB };
```

The previous example deletes all the rows in the table named TAB.

Note: For information on compiling and running SQLJ applications, see `Compile and run SQLJ programs`.

In an SQLJ process clause, the tokens that appear inside the braces are either SQL tokens or host variables. All host variables are distinguished by the colon (:) character. SQL tokens never occur outside the braces of an SQLJ process clause. For example, the following Java method inserts its arguments into an SQL table:

```
public void insertIntoTAB1 (int x, String y, float z) throws SQLException
{
    #sql { INSERT INTO TAB1 VALUES (:x, :y, :z) };
}
```

The method body consists of an SQLJ process clause containing the host variables x, y, and z. For more information on host variables, see `Host variables in SQLJ`.

In general, SQL tokens are case insensitive (except for identifiers delimited by double quotation marks), and can be written in upper, lower, or mixed case. Java tokens, however, are case sensitive. For clarity in

examples, case insensitive SQL tokens are uppercase, and Java tokens are lowercase or mixed case. Throughout this topic, the lowercase `null` is used to represent the Java "null" value, and the uppercase `NULL` is used to represent the SQL "null" value.

The following types of SQL constructs may appear in SQLJ programs:

- Queries
For example, `SELECT` statements and expressions.
- SQL Data Change statements (DML)
For example, `INSERT`, `UPDATE`, `DELETE`.
- Data statements
For example, `FETCH`, `SELECT..INTO`.
- Transaction Control statements
For example, `COMMIT`, `ROLLBACK`, etc.
- Data Definition Language (DDL, also known as Schema Manipulation Language) statements
For example, `CREATE`, `DROP`, `ALTER`.
- Calls to stored procedures
For example, `CALL MYPROC(:x, :y, :z)`
- Invocations of stored functions
For example, `VALUES(MYFUN(:x))`

For an example of embedded SQLJ, see Example: Embed SQL Statements in your Java application.

Host variables in Structured Query Language for Java: Arguments to embedded SQL statements are passed through host variables. Host variables are variables of the host language, and they can appear in SQL statements. Host variables have up to three parts:

- A colon (`:`) prefix.
- A Java^(TM) host variable that is a Java identifier for a parameter, variable, or field.
- An optional parameter mode identifier.

This mode identifier can be one of the following:

`IN`, `OUT`, or `INOUT`.

The evaluation of a Java identifier does not have side effects in a Java program, so it may appear multiple times in the Java code generated to replace an SQLJ clause.

The following query contains the host variable, `:x`. This host variable is the Java variable, field, or parameter `x` that is visible in the scope containing the query.

```
SELECT COL1, COL2 FROM TABLE1 WHERE :x > COL3
```

Example: Embed SQL Statements in your Java application: The following example SQLJ application, `App.sqlj`, uses static SQL to retrieve and update data from the `EMPLOYEE` table of the DB2 sample database.

Example: Embed SQL Statements in your Java^(TM) application:

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
import sqlj.runtime.*;
import sqlj.runtime.ref.*;

#sql iterator App_Cursor1 (String empno, String firstme) ; // 1
#sql iterator App_Cursor2 (String) ;
```

```

class App
{
    /**
     * Register Driver **
     */

    static
    {
        try
        {
            Class.forName("com.ibm.db2.jdbc.app.DB2Driver").newInstance();
        }
        catch (Exception e)
        {
            e.printStackTrace();
        }
    }

    /**
     * Main **
     */

    public static void main(String argv[])
    {
        try
        {
            App_Cursor1 cursor1;
            App_Cursor2 cursor2;

            String str1 = null;
            String str2 = null;
            long count1;

            // URL is jdbc:db2:dbname
            String url = "jdbc:db2:sample";

            DefaultContext ctx = DefaultContext.getDefaultContext();
            if (ctx == null)
            {
                try
                {
                    // connect with default id/password
                    Connection con = DriverManager.getConnection(url);
                    con.setAutoCommit(false);
                    ctx = new DefaultContext(con);
                }
                catch (SQLException e)
                {
                    System.out.println("Error: could not get a default context");
                    System.err.println(e);
                    System.exit(1);
                }
                DefaultContext.setDefaultContext(ctx);
            }
        }
    }
}

```



```

// retrieve data from the database
System.out.println("Retrieve some data from the database.");
#sql cursor1 = {SELECT empno, firstnme FROM employee}; // 2

// display the result set
// cursor1.next() returns false when there are no more rows
System.out.println("Received results:");
while (cursor1.next()) // 3
{
    str1 = cursor1.empno(); // 4
    str2 = cursor1.firstnme();

    System.out.print (" empno= " + str1);
    System.out.print (" firstname= " + str2);
    System.out.println("");
}
cursor1.close(); // 9

// retrieve number of employee from the database
#sql { SELECT count(*) into :count1 FROM employee }; // 5
if (1 == count1)
    System.out.println ("There is 1 row in employee table");
else
    System.out.println ("There are " + count1
        + " rows in employee table");

// update the database
System.out.println("Update the database.");
#sql { UPDATE employee SET firstnme = 'SHILI' WHERE empno = '000010' };

// retrieve the updated data from the database
System.out.println("Retrieve the updated data from the database.");
str1 = "000010";
#sql cursor2 = {SELECT firstnme FROM employee WHERE empno = :str1}; // 6

// display the result set
// cursor2.next() returns false when there are no more rows
System.out.println("Received results:");
while (true)
{
    #sql { FETCH :cursor2 INTO :str2 }; // 7
    if (cursor2.endFetch()) break; // 8

    System.out.print (" empno= " + str1);
    System.out.print (" firstname= " + str2);
    System.out.println("");
}
cursor2.close(); // 9

// rollback the update
System.out.println("Rollback the update.");
#sql { ROLLBACK work };
System.out.println("Rollback done.");
}
catch( Exception e )
{

```

```

        e.printStackTrace();
    }
}
}

```

1. **Declare iterators.** This section declares two types of iterators:

App_Cursor1

Declares column data types and names, and returns the values of the columns according to column name (Named binding to columns).

App_Cursor2

Declares column data types, and returns the values of the columns by column position (Positional binding to columns).

2. **Initialize the iterator.** The iterator object `cursor1` is initialized using the result of a query. The query stores the result in `cursor1`.
3. **Advance the iterator to the next row.** The `cursor1.next()` method returns a Boolean `false` if there are no more rows to retrieve.
4. **Move the data.** The named accessor method `empno()` returns the value of the column named `empno` on the current row. The named accessor method `firstnme()` returns the value of the column named `firstnme` on the current row.
5. **SELECT data into a host variable.** The SELECT statement passes the number of rows in the table into the host variable `count1`.
6. **Initialize the iterator.** The iterator object `cursor2` is initialized using the result of a query. The query stores the result in `cursor2`.
7. **Retrieve the data.** The FETCH statement returns the current value of the first column declared in the `ByPos` cursor from the result table into the host variable `str2`.
8. **Check the success of a FETCH..INTO statement.** The `endFetch()` method returns a Boolean `true` if the iterator is not positioned on a row, that is, if the last attempt to fetch a row failed. The `endFetch()` method returns `false` if the last attempt to fetch a row was successful. DB2 attempts to fetch a row when the `next()` method is called. A `FETCH...INTO` statement implicitly calls the `next()` method.
9. **Close the iterators.** The `close()` method releases any resources held by the iterators. You should explicitly close iterators to ensure that system resources are released in a timely fashion.

For background information on this example, see *Embed SQL Statements in your Java application*.

Compile and run SQLJ programs

If your Java^(TM) program has embedded SQLJ statements, you need to follow a special procedure to compile and run it.

Note: Before you begin, set up your CLASSPATH to contain the following:

- /QIBM/ProdData/Os400/Java400/ext/sqlj_classes.jar
- /QIBM/ProdData/Os400/Java400/ext/translator.zip
- /QIBM/ProdData/Os400/Java400/ext/runtime.zip

To compile and run your structured query language for Java (SQLJ) program, follow these steps:

1. Use the SQLJ translator, `sqlj`, on your Java source code with embedded SQL to generate Java source code and associated profiles. There is one profile generated for each connection.

For example, type in the following command:

```
sqlj MyClass.sqlj
```

where `MyClass.sqlj` is the name of your SQLJ file.

In this example, the SQLJ translator generates a `MyClass.java` source code file and any associated profiles. The associated profiles are named `MyClass_SJProfile0.ser`, `MyClass_SJProfile1.ser`, `MyClass_SJProfile2.ser`, and so on.

Note: The SQLJ translator automatically compiles the translated Java source code into a class file unless you explicitly turn off the compile option with the `-compile=false` clause.

2. Use the SQLJ Profile Customizer tool, `db2profrc`, to install DB2 SQLJ Customizers on generated profiles and create the DB2 packages on the local system.

For example, type in the command:

```
db2profrc MyClass_SJProfile0.ser
```

where `MyClass_SJProfile0.ser` is the name of the profile on which the DB2 SQLJ Customizer is run.

Note: This step is optional but is recommended to increase runtime performance.

3. Run the Java class file just like any other Java class file.

For example, type in the command:

```
java MyClass
```

where `MyClass` is the name of your Java class file.

Java SQL routines

➤ Your iSeries server provides the ability to access Java^(TM) programs from SQL statements and programs. This can be done using Java stored procedures and Java user-defined functions (UDFs). The iSeries server supports both the DB2 and SQLJ conventions for calling Java stored procedures and Java UDFs. Both Java stored procedures and Java UDFs can use Java classes that are stored in JAR files. The iSeries server uses stored procedures defined by the *SQLJ Part 1* standard to register JAR files with the database.

To access your Java applications from SQL statements and programs, see the following:

Use Java SQL routines

Take the following steps to use Java SQL routines:

- Write the Java methods for the routines.
- Compile the Java classes.
- Make the compiled classes accessible to the Java virtual machine used by the database.
- Register the routine with the database.
- Use the Java SQL procedure.

Java stored procedures

When using Java to write stored procedures, you can use the following possible parameter passing styles:

- JAVA parameter style
- DB2GENERAL parameter style

Java user-defined scalar functions

A Java scalar function returns one value from a Java program to the database. Like Java stored procedures, Java scalar functions use one of two parameter styles, JAVA and DB2GENERAL.

Java user-defined table functions

DB2 provides the ability for a function to return a table. This is useful for exposing information from outside the database to the database in table form.

SQLJ procedures that manipulate JAR files

Both Java stored procedures and Java UDFs can use Java classes that are stored in Java JAR files. Find the following information about SQLJ procedures that manipulate JAR files:

- SQLJ.INSTALL_JAR
- SQLJ.REMOVE_JAR
- SQLJ.REPLACE_JAR
- SQLJ.UPDATEJARINFO
- SQLJ.RECOVERJAR

Parameter passing conventions for Java stored procedures and UDFs

This section describes how SQL data types are represented in Java stored procedures and UDFs.



Use Java SQL routines

» You can access Java^(TM) programs from SQL statements and programs. This can be done using Java stored procedures and Java user-defined functions (UDFs).

To use Java SQL routines, the following steps must be taken:

1. Write the Java methods for the routines.

A Java SQL routine processes a Java method from SQL. This method must be written using either the DB2^(R) or SQLJ parameter passing conventions. See Java stored procedures, Java user-defined functions, and Java user-defined table functions for more information about coding a method used by a Java SQL routine.

2. Compile the Java classes.

Java SQL routines written using the Java parameter style may be compiled without any addition setup. However, Java SQL routines using the DB2GENERAL parameter style must extend either the `com.ibm.db2.app.UDF` class or `com.ibm.db2.app.StoredProc` class. These classes are contained in the JAR file, `/QIBM/ProdData/Java400/ext/db2routines_classes.jar`. When using `javac` to compile these routines, this JAR file must exist in the CLASSPATH. For example, the following command compiles a Java source file containing a routine which uses the DB2GENERAL parameter style:

```
javac -DCLASSPATH=/QIBM/ProdData/Java400/ext/db2routines_classes.jar
source.java
```

3. Make the compiled classes accessible to the JVM used by the database.

The user-defined classes used by the database Java virtual machine (JVM) can either reside in the `/QIBM/UserData/OS400/SQLLib/Function` directory or in a JAR file registered to the database.

The `/QIBM/UserData/OS400/SQLLib/Function` is the iSeries equivalent of `/sqllib/function`, the directory where DB2 UDB stores Java stored procedures and Java UDFs on other platforms. If the class is part of a Java package, it must reside in the appropriate subdirectory. For example, if the `runit` class is created as part of the `foo.bar` package, the file `runit.class` should be in the integrated file system directory, `/QIBM/ProdData/OS400/SQLLib/Function/foo/bar`.

The class file may also be placed in a JAR file that is registered to the database. The JAR file is registered using the `SQLJ.INSTALL_JAR` stored procedure. This stored procedure is used to assign a JAR ID to a JAR file. This JAR ID is used to identify the JAR file in which the class file resides. See SQLJ procedures that manipulate JAR files for more information on `SQLJ.INSTALL_JAR` as well as other stored procedures to manipulate JAR files.

4. Register the routine with the database.

Java SQL routines is registered with the database using the `CREATE PROCEDURE` and `CREATE FUNCTION` SQL statements. These statements contain the following elements:

CREATE keywords

The SQL statements to create a Java SQL routine begin with either `CREATE PROCEDURE` or `CREATE STATEMENT`.

Name of routine

The SQL statement then identifies the name of the routine that is known to the database. This is the name that is used to access the Java routine from SQL.

Parameters and return value

The SQL statement then identifies the parameters and return values, if applicable, for the Java routine.

LANGUAGE JAVA

The SQL statement uses the keywords LANGUAGE JAVA to indicate that the routine was written in Java.

PARAMETER STYLE KEYWORDS

The SQL statement then identifies the parameter style using the keywords PARAMETER STYLE JAVA or PARAMETER STYLE DB2GENERAL.

External name

The SQL statement then identifies the Java method to be processed as Java SQL routines. The external name has one of two formats:

- If the method is in a class file that is located under the /QIBM/UserData/OS400/SQLLib/Function directory, then the method is identified using the format *classname.methodname*, where *classname* is the fully qualified name of the class and *methodname* is the name of the method.
- If the method is in a JAR file registered to the database, then the method is identified using the format *jarid.classname.methodname*, where *jarid* is the JAR ID of the registered JAR file, *classname* is the name of the class, and *methodname* is the name of the method.

The iSeries Navigator may be used to create a stored procedure or user-defined function that uses the Java parameter style.

5. Use the Java procedure.

A Java stored procedure is called using the SQL CALL statement. A Java UDF is a function that is called as part of another SQL statement.

**Java stored procedures**

» When using Java^(TM) to write stored procedures, you can use two possible parameter passing styles. The recommended style is the JAVA parameter style, which matches the parameter style specified in the SQLJ: SQL routines standard. The second style, DB2GENERAL, is a parameter style defined by DB2^(R) UDB. The parameter style also determines the conventions that you must use when coding a Java stored procedure.

Additionally, you should also be aware of some restrictions that are placed on Java stored procedures. «

JAVA parameter style: » When you code a Java^(TM) stored procedure that uses the JAVA parameter style, you must use the following conventions:

- The Java method must be a public void static (not instance) method.
- The parameters of the Java method must be SQL-compatible types.
- A Java method may test for an SQL NULL value when the parameter is a null-capable type (like String).
- Output parameters are returned by using single element arrays.
- The Java method may access the current database using the getConnection method.

Java stored procedures using the JAVA parameter style are public static methods. Within the classes, the stored procedures are identified by their method name and signature. When you call a stored procedure, its signature is generated dynamically, based on the variable types defined by the CREATE PROCEDURE statement.

If a parameter is passed in a Java type that permits the null value, a Java method can compare the parameter to null to determine if an input parameter is an SQL NULL.

The following Java types do not support the null value:

- short
- int
- long
- float
- double

If a null value is passed to a Java type that does not support the null value, an SQL Exception with an error code of -20205 will be returned.

Output parameters are passed as arrays that contain one element. The Java stored procedure can set the first element of the array to set the output parameter.

A connection to the embedding application context is accessed using the following Java Database Connectivity (JDBC) call:

```
connection=DriverManager.getConnection("jdbc:default:connection");
```

This connection then runs SQL statements with JDBC APIs.

The following is a small stored procedure with one input and two outputs. It runs the given SQL query, and returns both the number of rows in the result and the SQLSTATE.

Example: Stored procedure with one input and two outputs

Note: Read the Code example disclaimer for important legal information.

```
package mystuff;

import java.sql.*;
public class sample2 {
    public static void donut(String query, int[] rowCount,
        String[] sqlstate) throws Exception {
        try {
            Connection c=DriverManager.getConnection("jdbc:default:connection");
            Statement s=c.createStatement();
            ResultSet r=s.executeQuery(query);
            int counter=0;
            while(r.next()){
                counter++;
            }
            r.close(); s.close();
            rowCount[0] = counter;
        }catch(SQLException x){
            sqlstate[0]= x.getSQLState();
        }
    }
}
```

In the SQLj standard, to return a result set in routines that use the JAVA parameter style, the result set must be set explicitly. When a procedure is created that returns result sets, additional result set parameters are added to the end of the parameter list. For example, the statement

```
CREATE PROCEDURE RETURN TWO()
DYNAMIC RESULT SETS 2
LANGUAGE JAVA
PARAMETER STYLE JAVA
EXTERNAL NAME 'javaClass!returnTwoResultSets'
```

would call a Java method with the signature `public static void returnTwoResultSets(ResultSet[] rs1, ResultSet[] rs2)`.

The output parameters of the result sets must be explicitly set as illustrated in the following example. As in the DB2GENERAL style, the result sets and corresponding statements should not be closed.

Example: Stored procedure that returns two result sets

Note: Read the Code example disclaimer for important legal information.

```
import java.sql.*;
public class javaClass {
    /**
     * Java stored procedure, with JAVA style parameters,
     * that processes two predefined sentences
     * and returns two result sets
     *
     * @param ResultSet[] rs1    first ResultSet
     * @param ResultSet[] rs2    second ResultSet
     */
    public static void returnTwoResultSets (ResultSet[] rs1, ResultSet[] rs2) throws Exception
    {
        //get caller's connection to the database; inherited from StoredProc
        Connection con = DriverManager.getConnection("jdbc:default:connection");

        //define and process the first select statement
        Statement stmt1 = con.createStatement();
        String sql1 = "select value from table01 where index=1";
        rs1[0] = stmt1.executeQuery(sql1);

        //define and process the second select statement
        Statement stmt2 = con.createStatement();
        String sql2 = "select value from table01 where index=2";
        rs2[0] = stmt2.executeQuery(sql2);
    }
}
```

On the server, the additional result set parameters are not examined to determine the ordering of the results sets. The results sets on the server are returned in the order in which they were opened. To ensure compatibility with the SQLj standard, the result should be assigned in the order that they are opened, as previously shown. <<

DB2GENERAL parameter style: >> When coding a Java^(TM) stored procedure that uses the DB2GENERAL parameter style, you must use the following conventions:

- The class that defines a Java stored procedure must *extend*, or be a subclass of, the Java `com.ibm.db2.app.StoredProc` class.
- The Java method must be a public void instance method.
- The parameters of the Java method must be SQL-compatible types.
- A Java method may test for a SQL NULL value using the `isNull` method.
- The Java method must explicitly set the return parameters using the `set` method.
- The Java method may access the current database using the `getConnection` method.

A class that includes a Java stored procedure must extend the class, `com.ibm.db2.app.StoredProc`. Java stored procedures are public instance methods. Within the classes, the stored procedures are identified by their method name and signature. When you call a stored procedure, its signature is generated dynamically, based on the variable types defined by the `CREATE PROCEDURE` statement.

The `com.ibm.db2.app.StoredProc` class provides the `isNull` method, which permits a Java method to determine if an input parameter is an SQL NULL. The `com.ibm.db2.app.StoredProc` class also provides

set...() methods that set output parameters. You must use these methods to set output parameters. If you do not set an output parameter, then the output parameter returns the SQL NULL value.

The com.ibm.db2.app.StoredProc class provides the following routine to fetch a JDBC connection to the embedding application context. A connection to the embedding application context is accessed using the following JDBC call:

```
public Java.sql.Connection getConnection( )
```

This connection then runs SQL statements with JDBC APIs.

The following is a small stored procedure with one input and two outputs. It processes the given SQL query, and returns both the number of rows in the result and the SQLSTATE.

Example: Stored procedure with one input and two outputs

Note: Read the Code example disclaimer for important legal information.

```
package mystuff;

import com.ibm.db2.app.*;
import java.sql.*;
public class sample2 extends StoredProc {
    public void donut(String query, int rowCount,
        String sqlstate) throws Exception {
        try {
            Statement s=getConnection().createStatement();
            ResultSet r=s.executeQuery(query);
            int counter=0;
            while(r.next()){
                counter++;
            }
            r.close(); s.close();
            set(2, counter);
        }catch(SQLException x){
            set(3, x.getSQLState());
        }
    }
}
```

To return a result set in procedures that use the DB2GENERAL parameter style, the result set and the responding statement must be left open at the end of the procedure. The result set that is returned must be closed by the client application. If multiple results sets are returned, they are returned in the order in which they were opened. For example, the following stored procedure returns two results sets.

Example: Stored procedure that returns two results sets

Note: Read the Code example disclaimer for important legal information.

```
public void returnTwoResultSets() throws Exception
{
    // get caller's connection to the database; inherited from StoredProc
    Connection con = getConnection ();
    Statement stmt1 = con.createStatement ();
    String sql1 = "select value from table01 where index=1";
    ResultSet rs1 = stmt1.executeQuery(sql1);
    Statement stmt2 = con.createStatement();
    String sql2 = "select value from table01 where index=2";
    ResultSet rs2 = stmt2.executeQuery(sql2);
}
```



Restrictions on Java stored procedures: » The following restrictions apply to Java[™] stored procedures:

- A Java stored procedure should not create additional threads. An additional thread may be created in a job only if the job is multithread capable. Because there is no guarantee that a job that calls an SQL stored procedure is multithread capable, a Java stored procedure should not create additional threads.
- You cannot use adopted authority to access Java class files.
- A Java stored procedure always uses the latest version of the Java Development Kit that is installed on the system.
- Since Blob and Clob classes reside in both the java.sql and com.ibm.db2.app packages, the programmer must use the entire name of these classes if both classes are used in the same program. The program must ensure that the Blob and Clob classes from the com.ibm.db2.app are used as the parameters passed to the stored procedure.
- When a Java stored procedure is created, the system generates a service program in the library. This service program is used to store the procedure definition. The service program has a name that is generated by the system. This name can be obtained by examining the job log of the job that created the stored procedure. If the service program object is saved and then restored, then the procedure definition is restored. If a Java stored procedure is to be moved from one system to another, you are responsible for moving the program that contains the procedure definition as well as the integrated file system file, which contains the Java class.
- A Java stored procedure cannot set the properties (for example, system naming) of the JDBC connection that is used to connect to the database. The default JDBC connection properties are always used, except when prefetching is disabled. «

Java user-defined scalar functions

» A Java[™] **scalar function** returns one value from a Java program to the database. For example, a scalar function could be created that returns the sum of two numbers. Like Java stored procedures, Java scalar functions use one of two parameter styles, Java and DB2GENERAL. When coding a Java user-defined function (UDF), you must be aware of restrictions placed on creating Java scalar functions.

Parameter style Java: The Java parameter style is the style specified by the *SQLJ Part 1: SQL Routines* standard. When coding a Java UDF, use the following conventions.

- The Java method must be a public static method.
- The Java method must return an SQL compatible type. The return value is the result of the method.
- The parameters of the Java method must be SQL compatible types.
- The Java method may test for a SQL NULL for Java types that permit the null value.

For example, given a UDF called sample!test3 that returns INTEGER and takes arguments of type CHAR(5), BLOB(10K), and DATE, DB2 expects the Java implementation of the UDF to have the following signature:

```
import com.ibm.db2.app.*;
public class sample {
    public static int test3(String arg1, Blob arg2, Date arg3) { ... }
}
```

The parameters of a Java method must be SQL compatible types. For example, if a UDF is declared as taking arguments of SQL types t1, t2, and t3, and returning type t4, it is called as a Java method with the expected Java signature:

```
public static T4 name (T1 a, T2 b, T3 c) { .....}
```

where:

- *name* is the method name
- T1 through T4 are the Java types that correspond to SQL types t1 through t4.
- *a*, *b*, and *c* are arbitrary variable names for the input arguments.

The correlation between SQL types and Java types is found in Parameter passing conventions for stored procedures and UDFs.

SQL NULL values are represented by Java variables that are not initialized. These variables have a Java null value if they are object types. If an SQL NULL is passed to a Java scalar data type, such as int, then an exception condition is raised.

To return a result from a Java UDF when using the JAVA parameter style, simply return the result from the method.

```
{ ....  
  return value;  
}
```

Like C modules used in UDFs and stored procedures, you cannot use the Java standard I/O streams (System.in, System.out, and System.err) in Java UDFs.

Parameter style DB2GENERAL: Parameter style DB2GENERAL is used by Java UDFs. In this parameter style, the return value is passed as the last parameter of the function and must be set using a *set* method of the com.ibm.db2.app.UDF class.

When coding a Java UDF, the following conventions must be followed:

- The class, which includes the Java UDF, must *extend*, or be a subclass of, the Java com.ibm.db2.app.UDF class.
- For the DB2GENERAL parameter style, the Java method must be a public void instance method.
- The parameters of the Java method must be SQL-compatible types.
- The Java method may test for an SQL NULL value using the isNull method.
- For the DB2GENERAL parameter style, the Java method must explicitly set the return parameter using the set() method.

A class that includes a Java UDF must extend the Java class, com.ibm.db2.app.UDF. A Java UDF that uses the DB2GENERAL parameter style must be a void instance method of the Java class. For example, for a UDF called sample!test3 that returns INTEGER and takes arguments of type CHAR(5), BLOB(10K), and DATE, DB2 expects the Java implementation of the UDF to have the following signature:

```
import com.ibm.db2.app.*;  
public class sample extends UDF {  
    public void test3(String arg1, Blob arg2, String arg3, int result) { ... }  
}
```

The parameters of a Java method must be SQL types. For example, if a UDF is declared as taking arguments of SQL types t1, t2, and t3, returning type t4, it is called as a Java method with the expected Java signature:

```
public void name (T1 a, T2 b, T3 c, T4 d) { .....}
```

where:

- *name* is the method name
- T1 through T4 are the Java types that correspond to SQL types t1 through t4.
- *a*, *b*, and *c* are arbitrary variable names for the input arguments.
- *d* is an arbitrary variable name that represents the UDF result being computed.

The correlation between SQL types and Java types is given in the section, Parameter passing conventions for stored procedures and UDFs.

SQL NULL values are represented by Java variables that are not initialized. These variables have a value of zero if they are primitive types, and Java null if they are object types, in accordance with Java rules. To tell an SQL NULL apart from an ordinary zero, the isNull method can be called for any input argument:

```
{ ....
  if (isNull(1)) { /* argument #1 was a SQL NULL */ }
  else           { /* not NULL */ }
}
```

In the previous example, the argument numbers start at one. The isNull() function, like the other functions that follow, are inherited from the com.ibm.db2.app.UDF class. To return a result from a Java UDF when using the DB2GENERAL parameter style, use the set() method in the UDF, as follows:

```
{ ....
  set(2, value);
}
```

Where 2 is the index of an output argument, and value is a literal or variable of a compatible type. The argument number is the index in the argument list of the selected output. In the first example in this section, the int result variable has an index of 4. An output argument that is not set before the UDF returns has a NULL value.



Like C modules used in UDFs and stored procedures, you cannot use the Java standard I/O streams (System.in, System.out, and System.err) in Java UDFs.

Typically, DB2 calls a UDF many times, once for each row of an input or result set in a query. If SCRATCHPAD is specified in the CREATE FUNCTION statement of the UDF, DB2 recognizes that some "continuity" is needed between successive invocations of the UDF, and therefore, for DB2GENERAL parameter style functions, the implementing Java class is not instantiated for each call, but generally speaking once per UDF reference per statement. If, however, NO SCRATCHPAD is specified for a UDF, then a clean instance is instantiated for each call to the UDF, by means of a call to the class constructor.

A scratchpad may be useful for saving information across calls to a UDF. Java UDFs can either use instance variables or set the scratchpad to achieve continuity between calls. Java UDFs access the scratchpad with the getScratchPad and setScratchPad methods available in com.ibm.db2.app.UDF. At the end of a query, if you specify the FINAL CALL option on the CREATE FUNCTION statement, the object's public void close() method is called (for DB2GENERAL parameter style functions). If you do not define this method, a stub function takes over and the event is ignored. The com.ibm.db2.app.UDF class contains useful variables and methods that you can use within a DB2GENERAL parameter style UDF. These variables and methods are explained in the following table.


Variables and Methods	Description
<pre>public static final int SQLUDF_FIRST_CALL = -1; public static final int SQLUDF_NORMAL_CALL = 0; ➤ public static final int SQLUDF_TF_FIRST = -2; public static final int SQLUDF_TF_OPEN = -1; public static final int SQLUDF_TF_FETCH = 0; public static final int SQLUDF_TF_CLOSE = 1; public static final int SQLUDF_TF_FINAL = 2; ⚡</pre>	<p>For scalar UDFs, these are constants to determine if the call is a first call or a normal call. For table UDFs, these are constants to determine if the call is a first call, open call, fetch call, close call, or final call.</p>
<pre>public Connection getConnection();</pre>	<p>The method obtains the JDBC connection handle for this stored procedure call and returns a JDBC object that represents the calling application's connection to the database. It is analogous to the result of a null SQLConnect() call in a C stored procedure.</p>


Variables and Methods	Description
public void close();	This method is called by the database at the end of a UDF evaluation, if the UDF was created with the FINAL CALL option. It is analogous to the final call for a C UDF. If a Java UDF class does not implement this method, this event is ignored.
public boolean isNull(int i)	This method tests whether an input argument with the given index is an SQL NULL.
public void set(int i, short s); public void set(int i, int j); public void set(int i, long j); public void set(int i, double d); public void set(int i, float f); public void set(int i, BigDecimal bigDecimal); public void set(int i, String string); public void set(int i, Blob blob); public void set(int i, Clob clob); public boolean needToSet(int i);	These methods set an output argument to the given value. An exception is thrown if anything goes wrong, including the following: <ul style="list-style-type: none"> • UDF call is not in progress • Index does not refer to valid output argument • Data type does not match • Data length does not match • Code page conversion error occurs
public void setSQLstate(String string);	This method may be called from a UDF to set the SQLSTATE to be returned from this call. If the string is not acceptable as an SQLSTATE, an exception is thrown. The user may set the SQLSTATE in the external program to return an error or warning from the function. In this case, the SQLSTATE must contain one of the following: <ul style="list-style-type: none"> • '00000' to indicate success • '01Hxx', where xx is any two digits or uppercase letters, to indicate a warning • '38yxx', where y is an uppercase letter between 'I' and 'Z' and xx is any two digits or uppercase letters, to indicate an error
public void setSQLmessage(String string);	This method is similar to the setSQLstate method. It sets the SQL message result. If the string is not acceptable (for example, longer than 70 characters), an exception is thrown.
public String getFunctionName();	This method returns the name of the processing UDF.
public String getSpecificName();	This method returns the specific name of the processing UDF.
public byte[] getDBinfo();	This method returns a raw, unprocessed DBINFO structure for the processing UDF, as a byte array. The UDF must have been registered (using CREATE FUNCTION) with the DBINFO option.
public String getDBname(); public String getDBauthid(); public String getDBver_rel(); public String getDBplatform(); public String getDBapplid(); public String getDBtblschema(); public String getDBtblname(); public String getDBcolname();	These methods return the value of the appropriate field from the DBINFO structure of the processing UDF. The UDF must have been registered (using CREATE FUNCTION) with the DBINFO option. The getDBtblschema(), getDBtblname(), and getDBcolname() methods only return meaningful information if a user-defined function is specified on the right side of a SET clause in an UPDATE statement.
public int getCCSID();	This method returns the CCSID of the job.
public byte[] getScratchpad();	This method returns a copy of the scratchpad of the currently processing UDF. You must first declare the UDF with the SCRATCHPAD option.

Variables and Methods	Description
<code>public void setScratchpad(byte ab[]);</code>	This method overwrites the scratchpad of the currently processing UDF with the contents of the given byte array. You must first declare the UDF with the SCRATCHPAD option. The byte array must have the same size as <code>getScratchpad()</code> returns.
<code>public int getCallType();</code>	This method returns the type of call that is currently being made. These values correspond to the C values defined in <code>sqludf.h</code> . Possible return values include the following: <ul style="list-style-type: none"> • SQLUDF_FIRST_CALL • SQLUDF_NORMAL_CALL •  SQLUDF_TF_FIRST • SQLUDF_TF_OPEN • SQLUDF_TF_FETCH • SQLUDF_TF_CLOSE • SQLUDF_TF_FINAL 



Restrictions on Java user-defined functions:  The following restrictions apply to Java^(TM) user-defined functions (UDFs):

- A Java UDF should not create additional threads. An additional thread may be created in a job only if the job is multithread capable. Since it cannot be guaranteed that a job that calls an SQL stored procedure is multithread capable, a Java stored procedure should not create additional threads.
- The complete name of the Java stored procedure defined to the database is limited to 279 characters. This limit is a consequence of the EXTERNAL_NAME column, which has a maximum width of 279 characters.
- Adopted authority cannot be used to access Java class files.
- A Java UDF always uses the latest version of the JDK that is installed on the system.
- Since Blob and Clob classes reside in both the `java.sql` and `com.ibm.db2.app` packages, the programmer must use the entire name of these classes if both classes are used in the same program. The program must ensure that the Blob and Clob classes from the `com.ibm.db2.app` are used as the parameters passed to the stored procedure.
- Like sourced functions, when a Java UDF is created, a service program in the library is used to store the function definition. The name of the service program is generated by the system and can be found in the job log of the job that created the function. If this object is saved and then restored to another system, then the function definition is restored. If a Java UDF is to be moved from one system to another, you are responsible for moving the service program that contains the function definition as well as the integrated file system file that contains the Java class.
- A Java UDF cannot set the properties (for example, system naming) of the JDBC connection that is used to connect to the database. The default JDBC connection properties are always used, except when prefetching is disabled. 

Java user-defined table functions:  DB2 provides the ability for a function to return a table. This is useful for exposing information from outside the database to the database in table form. For example, a table can be created that exposes the properties set in the Java^(TM) virtual machine (JVM) used for Java stored procedures and Java UDFs (both table and scalar).

The *SQLJ Part 1: SQL Routines* standard does support table functions. Consequently, table functions are only available using parameter style DB2GENERAL.

Five different types of calls are made to a table function. The following table explains these calls. These assume that scratchpad has been specified on the create function SQL statement.

Point in scan time	NO FINAL CALL LANGUAGE JAVA SCRATCHPAD	FINAL CALL LANGUAGE JAVA SCRATCHPAD
Before the first OPEN of the table function	No calls	Class constructor is called (means new scratchpad). UDF method is called with FIRST call.
At each OPEN of the table function.	Class constructor is called (means new scratchpad). UDF method is called with OPEN call.	UDF method is called with OPEN call.
At each FETCH for a new row of table function data.	UDF method is called with FETCH call.	UDF method is called with FETCH call.
At each CLOSE of the table function	UDF method is called with CLOSE call. The close() method, if it exists, is also called.	UDF method is called with CLOSE call.
After the last CLOSE of the table function.	No calls	UDF method is called with FINAL call. The close() method, if it exists, is also called.

Example: Java table function: The following is an example of a Java table function that determines the properties set in the JVM used to run the Java user-defined table function.

Note: Read the Code example disclaimer for important legal information.

```
import com.ibm.db2.app.*;
import java.util.*;

public class JVMProperties extends UDF {
    Enumeration propertyNames;
    Properties properties ;

    public void dump (String property, String value) throws Exception
    {
        int callType = getCallType();
        switch(callType) {
            case SQLUDF_TF_FIRST:
                break;
            case SQLUDF_TF_OPEN:
                properties = System.getProperties();
                propertyNames = properties.propertyNames();
                break;
            case SQLUDF_TF_FETCH:
                if (propertyNames.hasMoreElements()) {
                    property = (String) propertyNames.nextElement();
                    value = properties.getProperty(property);
                    set(1, property);
                    set(2, value);
                } else {
                    setSQLstate("02000");
                }
                break;
            case SQLUDF_TF_CLOSE:
                break;
            case SQLUDF_TF_FINAL:
                break;
            default:
                throw new Exception("UNEXPECT call type of "+callType);
        }
    }
}
```

After the table function is compiled, and its class file copied to /QIBM/UserData/OS400/SQLLib/Function, the function can be registered to the database by using the following SQL statement.

```
create function properties()  
returns table (property varchar(500), value varchar(500))  
external name 'JVMPProperties.dump' language java  
parameter style db2general fenced no sql  
disallow parallel scratchpad
```

After the function has been registered, it can be used as part of an SQL statement. For example, the following SELECT statement returns the table generated by the table function.


```
SELECT * FROM TABLE(PROPERTIES())
```




SQLJ procedures that manipulate JAR files

Both Java[™] stored procedures and Java UDFs can use Java classes that are stored in Java JAR files. To use a JAR file, a *jar-id* must be associated with the JAR file. The system provides stored procedures in the SQLJ schema that allow *jar-ids* and JAR files to be manipulated. These procedures allow JAR files to be installed, replaced, and removed. They also provide the ability to use and update the SQL catalogs associated with JAR files.

For more information, see the following topics:

- SQLJ.INSTALL_JAR
- SQLJ.REMOVE_JAR
- SQLJ.REPLACE_JAR
- SQLJ.UPDATEJARINFO
- SQLJ.RECOVERJAR

SQLJ.INSTALL_JAR:  The SQLJ.INSTALL_JAR stored procedure installs a JAR file into the database system. This JAR file can be used in subsequent CREATE FUNCTION and CREATE PROCEDURE statements.

Authorization: The privilege held by the authorization ID of the CALL statement must include at least one of the following for the SYSJAROBJECTS and SYSJARCONTENTS catalog tables:

- The following system authorities:
 - The INSERT and SELECT privileges on the table
 - The system authority *EXECUTE on library QSYS2
- Administrative authority

The privilege held by the authorization ID of the CALL statement must also have the following authorities:

- Read (*R) access to the JAR file specified in the *jar-url* parameter being installed.
- Write, Execute, and Read (*RWX) access to the directory where the JAR file is installed. This directory is /QIBM/UserData/OS400/SQLLib/Function/*jar/schema*, where *schema* is the schema of the *jar-id*.

Adopted authority cannot be used for these authorities.

SQL syntax:

```
>>>CALL--SQLJ.INSTALL_JAR-- (--'jar-url'--,--'jar-id'--,--deploy--)-->
```

```
>-----<
```

Description:

jar-url The URL containing the JAR file to be installed or replaced. The only URL scheme supported is 'file:'.

jar-id The JAR identifier in the database to be associated with the file specified by the *jar-url*. The *jar-id* uses SQL naming and the JAR file is installed in the schema or library specified by the implicit or explicit qualifier.

deploy

Value used to describe the *install_action* of the deployment descriptor file. If this integer is a nonzero value, then the *install_actions* of a deployment descriptor file should be run at the end of the *install_jar* procedure. The current version of DB2 UDB for iSeries only supports a value of zero.

Usage notes: When a JAR file is installed, DB2 UDB for iSeries registers the JAR file in the SYSJAROBJECTS system catalog. It also extracts the names of the Java^(TM) class files from the JAR file and registers each class in the SYSJARCONTENTS system catalog. DB2 UDB for iSeries copies the JAR file to a *jar/schema* subdirectory of the /QIBM/UserData/OS400/SQLLib/Function directory. DB2 UDB for iSeries gives the new copy of the JAR file the name given in the *jar-id* clause. A JAR file that has been installed by DB2 UDB for iSeries into a subdirectory of /QIBM/UserData/OS400/SQLLib/Function/*jar* should not be changed. Instead, the CALL SQLJ.REMOVE_JAR and CALL SQLJ.REPLACE_JAR SQL commands should be used to remove or replace an installed JAR file.

Example: The following command is issued from an SQL interactive session.

```
CALL SQLJ.INSTALL_JAR('file:/home/db2inst/classes/Proc.jar' , 'myproc_jar', 0)
```

The Proc.jar file located in the file:/home/db2inst/classes/ directory is installed into DB2 UDB for iSeries with the name myproc_jar. Subsequent SQL commands that use the Procedure.jar file refer to it with the name myproc_jar. <<

SQLJ.REMOVE_JAR: >> The SQLJ.REMOVE_JAR stored procedure removes a JAR file from the database system.

Authorization: The privilege held by the authorization ID of the CALL statement must include at least one of the following for the SYSJARCONTENTS and SYSJAROBJECTS catalog tables:

- The following system authorities:
 - The SELECT and DELETE privileges on the table
 - The system authority *EXECUTE on library QSYS2
- Administrative authority

The privilege held by the authorization ID of the CALL statement must also have the following authority.

- *OBJMGT authority to the JAR file being removed. The JAR file is named /QIBM/UserData/OS400/SQLLib/Function/*jar/schema/jarfile*.

Adopted authority cannot be used for this authority.

Syntax:

```
>>-CALL--SQLJ.REMOVE_JAR--(--'jar-id'--,--undeploy--)------><
```

Description:

jar-id The JAR identifier of the JAR file that is to be removed from the database.

undeploy

The value used to describe the *remove_action* of the deployment descriptor file. If this integer is a

non-zero value, then the `remove_actions` of a deployment descriptor file should be run at the end of the `install_jar` procedure. The current version of DB2 UDB for iSeries only supports a value of zero.

Example: The following command is issued from an SQL interactive session:

```
CALL SQLJ.REMOVE_JAR('myProc_jar', 0)
```

The JAR file `myProc_jar` is removed from the database. <<

SQLJ.REPLACE_JAR: >> The `SQLJ.REPLACE_JAR` stored procedure replaces a JAR file into the database system.

Authorization: The privilege held by the authorization ID of the `CALL` statement must include at least one of the following for the `SYSJAROBJECTS` and `SYSJARCONTENTS` catalog tables:

- The following system authorities:
 - The `SELECT`, `INSERT`, and `DELETE` privileges on the table
 - The system authority `*EXECUTE` on library `QSYS2`
- Administrative authority

The privilege held by the authorization ID of the `CALL` statement must also have the following authorities:

- Read (`*R`) access to the JAR file specified by the `jar-url` parameter being installed.
- `*OBJMGT` authority to the JAR file being removed. The JAR file is named `/QIBM/UserData/OS400/SQLLib/Function/jar/schema/jarfile`.

Adopted authority cannot be used for these authorities.

Syntax:

```
>>-CALL--SQLJ.REPLACE_JAR--(--'jar-url'--,--'jar-id'--)-----><
```

Description:

jar-url The URL containing the JAR file to be replaced. The only URL scheme supported is 'file:'.

jar-id The JAR identifier in the database to be associated with the file specified by the `jar-url`. The `jar-id` uses SQL naming and the JAR file is installed in the schema or library specified by the implicit or explicit qualifier.

Usage notes: The `SQLJ.REPLACE_JAR` stored procedure replaces a JAR file that was previously installed in the database using `SQLJ.INSTALL_JAR`.

Example: The following command is issued from an SQL interactive session:

```
CALL SQLJ.REPLACE_JAR('file:/home/db2inst/classes/Proc.jar' , 'myproc_jar')
```

The current JAR file referred to by the `jar-id` `myproc_jar` is replaced with the `Proc.jar` file located in the `file:/home/db2inst/classes/` directory. <<

SQLJ.UPDATEJARINFO: >> The `SQLJ.UPDATEJARINFO` updates the `CLASS_SOURCE` column of the `SYSJARCONTENTS` catalog table. This procedure is not part of the SQLJ standard but is used by the DB2 UDB for iSeries stored procedure builder.

Authorization: The privilege held by the authorization ID of the `CALL` statement must include at least one of the following for the `SYSJARCONTENTS` catalog table:

- The following system authorities:
 - The `SELECT` and `UPDATEINSERT` privileges on the table

- The system authority *EXECUTE on library QSYS2
- Administrative authority

The user running the CALL statement must also have the following authorities:

- Read (*R) access to the JAR file specified in the *jar-url* parameter. Read (*R) access to the JAR file being installed.
- Write, Execute, and Read (*RWX) access to the directory where the JAR file is installed. This directory is /QIBM/UserData/OS400/SQLLib/Function/jar/*schema*, where *schema* is the schema of the *jar-id*.

Adopted authority cannot be used for these authorities.

Syntax:

```
>>-CALL--SQLJ.UPDATEJARINFO--(--'jar-id'--,--'class-id'--,--'jar-url'--)-->
>-----<
```

Description:

jar-id The JAR identifier in the database that is to be updated.

class-id

The package qualified class name of the class to be updated.

jar-url The URL containing the classfile to update the JAR file with. The only URL scheme supported is 'file:'.

Example: The following command is issued from an SQL interactive session:

```
CALL SQLJ.UPDATEJARINFO('myproc_jar', 'mypackage.myclass',
                        'file:/home/user/mypackage/myclass.class')
```

The JAR file associated with the *jar-id* myproc_jar, is updated with a new version of the mypackage.myclass class. The new version of the class is obtained from the file /home/user/mypackage/myclass.class. <<

SQLJ.RECOVERJAR: >> The SQLJ.RECOVERJAR procedure takes the JAR file that is stored in the SYSJAROBJECTS catalog and restores it to the /QIBM/UserData/OS400/SQLLib/Function/jar/*schema*/*jar_id*.jar file.

Authorization: The privilege held by the authorization ID of the CALL statement must include at least one of the following for the SYSJAROBJECTS catalog table:

- The following system authorities:
 - The SELECT and UPDATEINSERT privileges on the table
 - The system authority *EXECUTE on library QSYS2
- Administrative authority

The user running the CALL statement must also have the following authorities:

- Write, Execute, and Read (*RWX) access to the directory where the JAR file is installed. This directory is /QIBM/UserData/OS400/SQLLib/Function/jar/*schema*, where *schema* is the schema of the *jar-id*.
- *OBJMGT authority to the JAR file being removed. The JAR file is named /QIBM/UserData/OS400/SQLLib/Function/jar/*schema*/jarfile.

Syntax:

```
>>-CALL--SQLJ.RECOVERJAR--(--'jar-id'--)-----<
```

Description:

jar-id The JAR identifier in the database that is to be recovered.

Example: The following command is issued from a SQL interactive session:

```
CALL SQLJ.UPDATEJARINFO('myproc_jar')
```

The JAR file associated with the myproc_jar is updated with the contents from SYSJARCONTENT table. The file is copied to /QIBM/UserData/OS400/SQLLib/Function/jar/jar_schema myproc_jar.jar. <<

Parameter passing conventions for Java stored procedures and UDFs

>> The following table lists how SQL data types are represented in Java^(TM) stored procedures and UDFs.

SQL data type	Java parameter style JAVA	Java parameter style DB2GENERAL
SMALLINT	short	short
INTEGER	int	int
BIGINT	long	long
DECIMAL(p,s)	BigDecimal	BigDecimal
NUMERIC(p,s)	BigDecimal	BigDecimal
REAL or FLOAT(p)	float	float
DOUBLE PRECISION or FLOAT or FLOAT(p)	double	double
CHARACTER(n)	String	String
VARCHAR(n)	String	String
VARCHAR(n) FOR BIT DATA	byte[]	com.ibm.db2.app.Blob
GRAPHIC(n)	String	String
VARGRAPHIC(n)	String	String
DATE	Date	String
TIME	Time	String
TIMESTAMP	Timestamp	String
Indicator Variable	-	-
CLOB	-	com.ibm.db2.app.Clob
BLOB	-	com.ibm.db2.app.Blob
DBCLOB	-	com.ibm.db2.app.Clob
DataLink	-	-



Java with other programming languages

With Java^(TM), you have multiple ways to call code that was written in languages other than Java.

Java Native Interface

One of the ways you can call code written in another language is to implement selected Java methods as 'native methods.' Native methods are procedures, written in another language, that provide the actual implementation of a Java method. Native methods can access the Java virtual machine using the Java Native Interface (JNI). These native methods run under the Java thread, which is a kernel thread, so they must be thread safe. A function is thread safe if you can start it simultaneously in multiple threads within the same process. A function is thread safe if and only if all the functions it calls are also thread safe.

Native methods are a "bridge" to access system functions that are not directly supported in Java, or to interface to existing user code. Use caution when using native methods, because the code that is being called may not be thread safe. >> See Use the Java Native Interface for native methods for more information about JNI and ILE native methods. <<

>> OS/400 PASE native methods

The iSeries Java virtual machine (JVM) now supports the use of native methods running in the OS/400^(R) PASE environment. OS/400 PASE native methods for Java enables you to easily port your Java applications that run in AIX^(R) to your iSeries server. You can copy the class files and AIX native method libraries to the integrated file system on the iSeries and run them from any of the control language (CL), Qshell or OS/400 PASE terminal session command prompts. <<

java.lang.Runtime.exec()

You can use `java.lang.Runtime.exec()` to call programs or commands from within a Java program. The `exec()` method starts another process in which any iSeries program or command can run. In this model, you can use standard in, standard out, and standard err of the child process for interprocess communication.

Interprocess communication

One option is to use sockets for interprocess communication between the parent and child processes.

You can also use stream files for communication between programs. Or see interprocess communication examples for an overview of your options when communicating with programs that are running in another process.

>> To call Java from other languages, see Example: Call Java from C or Example: Call Java from RPG for more information. <<

You can also use the IBM Toolbox for Java to call existing programs and commands on the iSeries server. Data queues and iSeries messages are usually used for interprocess communication with the IBM Toolbox for Java.

Note: By using `Runtime.exec()`, IBM Toolbox for Java, or JNI, you may compromise the portability of the Java program. You should avoid using these methods in a "pure" Java environment.

Java Invocation API


Using the Java Invocation API, which is also a part of the Java Native Interface (JNI) specification, allows a non-Java application to use the Java virtual machine. It also allows the use of Java code as an extension of the application.

Use the Java Native Interface for native methods

You should only use native methods in cases where pure Java^(TM) cannot meet your programming needs. Limit the use of native methods by only using them under these circumstances:

- To access system functions that are not available using pure Java.
- To implement extremely performance-sensitive methods that can benefit significantly from a native implementation.
- To interface to existing application program interfaces (API) that allow Java to call other APIs.

To use the Java Native Interface (JNI) for native methods, do these steps:

1. Design the class by specifying which methods are native methods with the standard Java language syntax.
2. Decide on a library and program name for the service program (*SRVPGM) that contains native method implementations. When coding the System.loadLibrary() method call in the static initializer for the class, specify the name of the service program.
3. Use the javac tool to compile the Java source into a class file.
4. Use the javah tool to create the header file (.h). This header file contains the exact prototypes for creating the native method implementations. The -d option specifies the directory where you should create the header file.
5. Copy the header file from the integrated file system into a member in a source file by using the Copy From Stream File (CPYFRMSTMF) command. You must copy the header file into a source file member for the C compiler to use it. Use the new stream file support for the Create Bound ILE C/400 Program (CRTCMOD) command to leave your C source and C header files in the integrated file system. For more information on the CRTCMOD command and the use of stream files, see the WebSphere Development Studio: ILE C/C++ Programmer's Guide, SC09-2712  .
6. Write the native method code. See Java native methods and threads considerations for details about the languages and functions that are used for native methods.
 - a. Include the header file that was created in the previous steps.
 - b. Match the prototypes in the header file exactly.
 - c. Convert strings to American Standard Code for Information Interchange (ASCII) if the strings are to pass to the Java virtual machine. For more information, see Java character encodings.
7. If your native method must interact with the Java virtual machine, use the functions that are provided with JNI.
8. Compile your C source code, using the CRTCMOD, into a module (*MOD) object.
9. Bind one or more module objects into a service program (*SRVPGM) by using the Create Service Program (CRTSRVPGM) command. The name of this service program must match the name that you supplied in your Java code that is in the System.load() or System.loadLibrary() function calls.
10. If you used the System.loadLibrary() call in your Java code, do one of the following.

If you are using a version that is older than J2SDK:

Add the iSeries library that contains your new service program to the iSeries library list. To add the library, use the Add Library List Entry (ADDLIBLE) command. This allows the Java program to find the service program when it processes the System.loadLibrary() function.

If you are using J2SDK, version 1.2 or higher:

You do not need to change your library list. Instead, you can either:

- Include the list of the libraries that you need in the LIBPATH environment variable. You can change the LIBPATH environment variable in QShell and from the iSeries command line.
 - From the Qshell command prompt, type in:


```
export LIBPATH=/QSYS.LIB/MYLIB.LIB
>> java -Djava.version=1.4 myclass<<
```
 - Or, from the command line:


```
ADDENVVAR LIBPATH '/QSYS.LIB/MYLIB.LIB'
>> JAVA PROP((java.version 1.4)) myclass<<
```
- Or, supply the list in the **java.library.path** property. You can change the java.library.path property in QShell and from the iSeries command line.
 - From the Qshell command prompt, type in:


```
>> java -Djava.library.path=/QSYS.LIB/MYLIB.LIB -Djava.version=1.4 myclass<<
```
 - Or, from the iSeries command line, type in:


```
>> JAVA PROP((java.library.path '/QSYS.LIB/MYLIB.LIB') (java.version '1.4')) myclass<<
```


Where */QSYS.LIB/MYLIB.LIB* is the library that you want to load using the `System.loadLibrary()` call, and *myclass* is the name of your Java application.

11. The patches syntax for `System.load(String patches)` can be any of these:
 - "path" (integrated file system filename that specifies the library that the service program provides), which is a symbolic link to a *SRVPGM, such as `"/qsys.lib/mylib.lib/myNMsp.srvpgm"`
 - `/qsys.lib/sysNMsp.srvpgm`
 - `/qsys.lib/mylib.lib/myNMsp.srvpgm`
 - If you have a version older than J2SDK: `/qsys.lib/%lib%.lib/myNMsp.srvpgm`
Note: This is equivalent to using the `System.loadLibrary("myNMsp")` method.

Note: If the pathname is used as a string literal, then you must enclose it in quotation marks. For example, `System.load("/qsys.lib/mylib.lib/myNMsp.srvpgm")`.

12. The "libya" syntax for `System.loadLibrary(String libya)` is `mypsp`. The system finds `mypsp` by using *LIBL. For example, `loadLibrary("myNMsp")` is equivalent to `System.load("/qsys.lib/%lib%.lib/myNMsp.srvpgm")`. If the "pathname" is used as a string literal, then you must enclose the libname in quotation marks.

Note: The `%lib%.lib` syntax is not supported for J2SDK.

For a complete description of the JNI, refer to the Java Native Interface by Sun Microsystems, Inc., and The Source for Java Technology java.sun.com .

See Examples: Use the Java Native Interface for native methods for an example of how to use the JNI for native methods.

Java Invocation API

The Invocation API, which is part of the Java^(TM) Native Interface (JNI), allows non-Java code to create a Java virtual machine, and load and use Java classes. This function lets a multithreaded program make use of Java classes that are running in a single Java virtual machine in multiple threads.

The application controls the Java virtual machine. The application can create the Java virtual machine, call Java methods (similar to the way in which an application calls subroutines), and destroy the Java virtual machine. Once you create the Java virtual machine, it remains ready to run within the process until the application explicitly destroys it. While being destroyed, the Java virtual machine performs clean-up, such as running finalizers, ending Java virtual machine threads, and releasing Java virtual machine resources.

With a Java virtual machine that is ready to run, an application written in C can call into the Java virtual machine to perform any function. It also can return from the Java virtual machine to the C application, call into the Java virtual machine again, and so on. The Java virtual machine is created once and does not have to be re-created before calling into the Java virtual machine to run a little or a lot of Java code.

When using the Invocation API to run Java programs, the destination for `STDOUT` and `STDERR` is controlled by the use of an environment variable called `QIBM_USE_DESCRIPTOR_STDIO`. If this environment variable is set to `Y` or `I` (for example, `QIBM_USE_DESCRIPTOR_STDIO=Y`), the Java virtual machine uses file descriptors for `STDIN` (fd 0), `STDOUT` (fd 1), and `STDERR` (fd 2). In this case, the program must set these file descriptors to valid values by opening them as the first three files or pipes in this job. The first file opened in the job is given fd of 0, the second fd of 1, and third is fd of 2. For jobs initiated with the spawn API, these descriptors can be preassigned using a file descriptor map (see documentation on Spawn API). If the environment variable `QIBM_USE_DESCRIPTOR_STDIO` is not set or is set to any other value, file descriptors are not used for `STDIN`, `STDOUT`, or `STDERR`. Instead, `STDOUT` and `STDERR` are routed to a spooled file that is owned by the current job, and use of `STDIN` results in an IO exception.

For an example that uses the Invocation API, see Example: Java Invocation API. See Invocation API functions for details about the Invocation API functions that are supported by the IBM Developer Kit for Java.

Invocation API functions: The IBM Developer Kit for Java^(TM) supports these Invocation API functions.

Note: Before using this API, you must ensure that you are in a multithread-capable job. See Multithreaded applications for more information about multithread-capable jobs.

- **JNI_GetDefaultJavaVMInitArgs**

Note: This function is only supported for Java Development Kit (JDK) 1.1.x.

Returns a JDK 1.1 structure that contains default values for arguments that need to be passed to `JNI_CreateJavaVM` when you create a Java virtual machine.

Signature:

```
jint JNI_GetDefaultJavaVMInitArgs(void *args_);
```

- **JNI_GetCreatedJavaVMs**

Returns information about all Java virtual machines that were created.

Signature:

```
jint JNI_GetCreatedJavaVMs(JavaVM **vmBuf,  
                           jsize bufLen,  
                           jsize *nVMs);
```

`vmBuf` is an output area whose size is determined by `bufLen`, which is the number of pointers. Each Java virtual machine has an associated `JavaVM` structure that is defined in `java.h`. This API stores a pointer to the `JavaVM` structure that is associated with each created Java virtual machine into `vmBuf`, unless `vmBuf` is full. Pointers to `JavaVM` structures are stored in the order of the corresponding Java virtual machines that are created. `nVMs` returns the number of virtual machines that are currently created. Your iSeries server supports the creation of more than one Java virtual machine, so you may expect a value higher than one. This information, along with the size of the `vmBuf`, determines whether pointers to `JavaVM` structures for each created Java virtual machine are returned.

- **JNI_CreateJavaVM**

Allows you to create a Java virtual machine and subsequently use it in an application.

Signature for Java Development Kit 1.1.x:

```
jint JNI_CreateJavaVM(JavaVM **p_vm,  
                     JNIEnv **p_env,  
                     void *vm_args);
```

Signature for Java 2 Software Development Kit (J2SDK):

```
jint JNI_CreateJavaVM(JavaVM **p_vm,  
                     void **p_env,  
                     void *vm_args);
```

`p_vm` is the address of a `JavaVM` pointer for the newly created Java virtual machine. Several other JNI Invocation APIs use `p_vm` to identify the Java virtual machine. `p_env` is the address of a JNI Environment pointer for the newly created Java virtual machine. It points to a table of JNI functions that start those functions. `vm_args` is a structure that contains Java virtual machine initialization parameters. When using JDK 1.1.x, you can obtain a structure that contains default values by calling `JNI_GetDefaultJavaVMInitArgs`. For details on how to do this with J2SDK, see Java Native Interface



If you start a Run Java (RUNJAVA) command or JAVA command and specify a property that has an equivalent command parameter, then the command parameter takes precedence. The property is ignored. For example, the `os400.optimization` parameter is ignored in this command:

```
JAVA CLASS>Hello) PROP((os400.optimization 0))
```

For a list of OS/400 unique properties that are supported by the JNI_CreateJavaVM API, see Java system properties.

Note:With multiple Java virtual machines within one process, all Java virtual machines share the same process static storage that is allocated for any native methods. Java virtual machine internal implementation already partitions data on a per-Java virtual machine basis, but you must consider that with native method applications, Java virtual machines share process static storage. For other considerations, see Support for multiple Java virtual machines.

- **DestroyJavaVM**

Destroys the Java virtual machine.

Signature:

```
jint DestroyJavaVM(JavaVM *vm)
```

When the Java virtual machine is created, vm is the JavaVM pointer that is returned.

- **AttachCurrentThread**

Attaches a thread to a Java virtual machine, so it can use Java virtual machine services.

Signature for Java Development Kit (JDK) 1.1.x:

```
jint AttachCurrentThread(JavaVM *vm,
                        JNIEnv **p_env,
                        void *thr_args);
```

Signature for Java 2 Software Development Kit (J2SDK):

```
jint AttachCurrentThread(JavaVM *vm,
                        void **p_env,
                        void *thr_args);
```

The JavaVM pointer, vm, identifies the Java virtual machine to which the thread is being attached. p_env is the pointer to the location where the JNI Interface pointer of the current thread is placed. thr_args contains VM specific thread attachment arguments.

- **DetachCurrentThread**



Signature:

```
jint DetachCurrentThread(JavaVM *vm);
```

vm identifies the Java virtual machine from which the thread is being detached.

For a complete description of the Invocation API functions, refer to the Java Native Interface Specification by Sun Microsystems, Inc., or The Source for Java Technology java.sun.com .

Support for multiple Java virtual machines: Java^(TM) on the iSeries server, unlike the Sun Microsystems, Inc. reference implementation, supports the creation of multiple Java virtual machines within a single job or process. This means that you can successfully call JNI_CreateJavaVM() more than once in a job, and JNI_GetCreatedJavaVMs() can return more than one Java virtual machine in its list of results.

-  Prior to V5R2, it would be possible for the JNI_GetCreatedJavaVMs function to return more than one JVM in its list of JVMs. In V5R2, the JNI_GetCreatedJavaVMs will return, at most, one JVM.
- Prior to V5R2, the JNI_CreateJavaVM function could be called repeatedly within a single process, and with each successful invocation, a separate and distinct JVM was created. In V5R2, the JNI_CreateJavaVM function returns an error code. (JNI_ERR -1 error defined in jni.h) 

If you want to create multiple Java virtual machines for use within a single job or process, you should carefully consider the following:

Native method static storage scoping

- Service programs that contain native method implementations are only activated once per job, regardless of the number of Java virtual machines that you create. This implies that native method static storage is scoped to the job and not to any specific Java virtual machine.
- Values that a native method places in static storage are independent of the Java virtual machine that called the native method. These values are visible to any Java virtual machine in the job.
- If you intentionally use native method static storage in a multiple Java virtual machine scenario, you should carefully consider the possible requirements for synchronization, above and beyond the use of synchronized methods and monitors, which are Java virtual machine specific. The qualification of a native method as synchronized only prevents simultaneous runs within a single Java virtual machine and not simultaneous runs from multiple Java virtual machines.

Java virtual machine stoppage

- If a Java virtual machine is abnormally stopped, either due to a user calling `java.lang.System.exit()` or an internal Java virtual machine failure, the failing Java virtual machine and all of its attached threads are stopped.
- If the process initial thread is among the threads that were attached to the failed Java virtual machine, an exception is thrown to the initial thread. If the initial thread handles this exception, the other Java virtual machines can continue to run.
- If the process initial thread is stopped, either by an unhandled exception or any other reason, all of the Java virtual machines in the process are stopped as well.

Abnormal stoppage from C

If you use the ILE/C `exit()` or `abort()` routines in any thread of a multithreaded job, you immediately bring down the entire job, including all Java virtual machines.

Example: Java Invocation API: This example follows the standard Invocation API paradigm. For example, it does the following:

- Creates a Java^(TM) virtual machine by using `JNI_CreateJavaVM`.
- Uses the Java virtual machine to find the class file that you want to run.
- Finds the methodID for the main method of the class.
- Calls the main method of the class.
- Reports errors if an exception occurs.

To compile this program, you must bind it with a service program that exports the functions to start a new Java virtual machine. These are the entry points that you need:

- `JNI_GetDefaultJavaVMInitArgs`, which initializes the parameters to create.
- `JNI_CreateJavaVM`, which creates the Java virtual machine.

When you compile the program, you do not need to do anything explicit with the compile command. The service program that exports these entry points is in the system binding directory. The name of the service program is `QJVAJNI`.

To run this program, use `SBMJOB CMD(CALL PGM(YOURLIB/PGMNAME)) ALWMLTTHD(*YES)`. Any job that creates a Java virtual machine must be multithread-capable. The only job on an iSeries server that is multithread-capable is a batch immediate (BCI) job. The output from the main program, as well as any output from the program, ends up in QPRINT spooled files. These spooled files are visible if you use the Work with Submitted Jobs (`WRKSBMJOB`) command and view the job that you started with Submit Job (`SBMJOB`).

Note: The C runtime `exit()` routine used below is not recommended unless you know that your program is the only thread in the process. When called from a process that is capable of supporting multiple threads, `exit()` immediately ends all threads in a process.

Example: Use the Java Invocation API with JDK 1.1.x.

Note: Read the Code example disclaimer for important legal information.

```
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <jni.h>

int main (int argc, char *argv[])
{
    JDK1_1InitArgs initArgs; /* Virtual Machine (VM) initialization structure.
        * This is the structure that is passed by reference to JNI_CreateJavaVM().
        * See jni.h for details.
        */

    JavaVM* myJVM;          /* The JavaVM and JNIEnv pointers that you get back. */

    JNIEnv* myEnv;         /* ...from the JNI_CreateJavaVM() call. */

    char*   myClasspath;   /* You need to change the classpath, so you have your own. */

    jclass myClass;        /* The class you are going to find, called 'NativeHello'. */

    jmethodID mainID;      /* The method ID of the class' "main" routine. */

    jclass  stringClass;   /* Required to create a string array argument for 'main'. */

    jobjectArray args;     /* Because main expects an array of strings, you must pass one. */

    /* Set the version field of the initialization arguments. */
    initArgs.version = 0x00010001;

    /* Get the default initialization arguments. */
    JNI_GetDefaultJavaVMInitArgs(&initArgs);

    /* Now, you want to add the directory onto the end of the classpath,
     * so that the findClass finds it correctly. To do this, you have two options:
     * You can append your classpath entries to the default classpath that is returned
     * by the call to JNI_GetDefaultJavaVMInitArgs, or
     * you can use OS/400 specific functions for the same result. This is a
     * three-step solution:
     * 1. Set the CLASSPATH environment variable to its requirements with 'putenv()'
     * 2. Clear the initialization arguments classpath to NULL, which forces
     *    JNI_CreateJavaVM to look at the CLASSPATH value
     * 3. Set the "os400.class.path.system=PRE" property, to force JNI_CreateJavaVM
     *    to prepend the system default classpath to the effective classpath.
     *
     * The first option is used in this example, because it is more platform independent
     *
     * Note: You must specify the directory name in UTF-8 format! So, you wrap
     * blocks of code in #pragma convert statements.
     */

    #pragma convert(819)
    myClasspath = malloc( strlen(initArgs.classpath) + strlen(":/CrtJvmExample") + 1 );
    strcpy( myClasspath, initArgs.classpath );
    strcat( myClasspath, ":/CrtJvmExample" );
    initArgs.classpath = myClasspath;

    #pragma convert(0)

    /* Create the JVM. */
    if (JNI_CreateJavaVM(&myJVM, &myEnv, &initArgs)) {
        fprintf(stderr, "Failed to create the JVM\n");
        exit(1);
    }
}
```

```

/* Use the newly created JVM to find the example class.
 * Note: Again, you are dealing with UTF-8 here, so you
 * have to wrap the calls in #pragma convert.
 */

#pragma convert(819)
if (! (myClass = (*myEnv)->FindClass(myEnv, "NativeHello"))) {

#pragma convert(0)
/* Cannot find the class, so write an error message
 * to C stderr and exit the program.
 */

fprintf(stderr, "Failed to find the class 'NativeHello'\n");
exit(1); /* Exit stops the entire process on an iSeries server. */
}

/* Now, get the method identifier for the 'main' entry point
 * of the class. Note: The signature of 'main' is always
 * the same for every class, "main" and "([Ljava/lang/String;)V"
 * Again, you are dealing with UTF-8.
 */

#pragma convert(819)
if (! (mainID = (*myEnv)->GetStaticMethodID(myEnv, myClass,
                                           "main",
                                           "([Ljava/lang/String;)V"))) {
/* The 'main' methodID is not found for some reason. */
if ( (*myEnv)->ExceptionOccurred(myEnv) ) {
/* a java exception occurred, so print it out */
(*myEnv)->ExceptionDescribe(myEnv);
/* The JVM ends. */
(*myEnv)->FatalError(myEnv, "Failed to find jmethodID of 'main()'");
}
}

#pragma convert(0)
/* Cannot find the 'main' methodID, so write an error message
 * to C stderr and exit the program.
 */

fprintf(stderr, "Failed to find the 'main()' method\n");
exit(1); /* Exit stops the entire process on an iSeries server. */
}

#pragma convert(819)
if (! (stringClass = (*myEnv)->FindClass(myEnv,"java/lang/String"))) {

#pragma convert(0)
/* Did not find java/lang/String, so write an error message
 * to C stderr and exit the program.
 */

fprintf(stderr, "Failed to find the java/lang/String");
exit(1); /* exit stops the entire process on an iSeries server.*/
}

/* Now, you need to create an empty array of strings,
 * because ([Ljava/lang/String) is a required part of the signature of
 * every Java main routine.
 */

if (! (args = (*myEnv)->NewObjectArray(myEnv,0,stringClass,0))) {
/* Empty array was not created, so write an error message
 * to C stderr and exit the program.
 */

```

```

        fprintf(stderr, "Failed to create empty array of strings");
        exit(1); /* Exit stops the entire process on an iSeries server. */
    }

/* Now, you have the methodID of main, and the class, so you can call the main method. */

    (*myEnv)->CallStaticVoidMethod(myEnv,myClass,mainID,args);

/* Check for errors. */
    if ( (*myEnv)->ExceptionOccurred(myEnv) ) {
        fprintf(stderr,"An exception occurred while running 'main'");
        exit(1);
    }

/* Finally, destroy the JavaVM that you created. */

    if ( (*myJVM)->DestroyJavaVM(myJVM) ) {
        fprintf(stderr, "Failed to destroy the JVM\n");
        exit(1);
    }

/* All done. */

    return 0;
}

```



Example: Use the Java Invocation API with J2SDK.

Note: Read the Code example disclaimer for important legal information.

```

#define OS400_JVM_12
#include <stdlib.h>
#include <stdio.h>
#include <fcntl.h>
#include <string.h>
#include <jni.h>

/* Specify the pragma that causes all literal strings in the
 * source code to be stored in ASCII (which, for the strings
 * used, is equivalent to UTF-8)
 */

#pragma convert(819)

/* Procedure: Oops
 *
 * Description: Helper routine that is called when a JNI function
 *             returns a zero value, indicating a serious error.
 *             This routine reports the exception to stderr and
 *             ends the JVM abruptly with a call to FatalError.
 *
 * Parameters:  env -- JNIEnv* to use for JNI calls
 *             msg -- char* pointing to error description in UTF-8
 *
 * Note:       Control does not return after the call to FatalError
 *             and it does not return from this procedure.
 */

void Oops(JNIEnv* env, char *msg) {
    if ((*env)->ExceptionOccurred(env)) {
        (*env)->ExceptionDescribe(env);
    }
    (*env)->FatalError(env, msg);
}

```

```

/* This is the program's "main" routine. */
int main (int argc, char *argv[])
{

    JavaVMInitArgs initArgs; /* Virtual Machine (VM) initialization structure, passed by
    * reference to JNI_CreateJavaVM(). See jni.h for details
    */

    JavaVM* myJVM;          /* JavaVM pointer set by call to JNI_CreateJavaVM */
    JNIEnv* myEnv;         /* JNIEnv pointer set by call to JNI_CreateJavaVM */
    char* myClasspath;     /* Changeable classpath 'string' */
    jclass myClass;        /* The class to call, 'NativeHello'. */
    jmethodID mainID;      /* The method ID of its 'main' routine. */
    jclass stringClass;    /* Needed to create the String[] arg for main */
    jobjectArray args;     /* The String[] itself */
    JavaVMOption options[1]; /* Options array -- use options to set classpath */
    int fd0, fd1, fd2;     /* file descriptors for IO */

    /* Open the file descriptors so that IO works. */
    fd0 = open("/dev/null1", O_CREAT|O_TRUNC|O_RDWR, S_IRUSR|S_IROTH);
    fd1 = open("/dev/null12", O_CREAT|O_TRUNC|O_WRONLY, S_IWUSR|S_IWOTH);
    fd2 = open("/dev/null13", O_CREAT|O_TRUNC|O_WRONLY, S_IWUSR|S_IWOTH);

    /* Set the version field of the initialization arguments for J2SDK. */
    initArgs.version = 0x00010002;

    /* Now, you want to specify the directory for the class to run in the classpath.
    * with Java2, classpath is passed in as an option.
    * Note: You must specify the directory name in UTF-8 format. So, you wrap
    * blocks of code in #pragma convert statements.
    */
    options[0].optionString="-Djava.class.path=/CrtJvmExample";

    initArgs.options=options; /* Pass in the classpath that has been set up. */
    initArgs.nOptions = 1; /* Only passing the one classpath option */

    /* Create the JVM -- a nonzero return code indicates there was
    * an error. Drop back into EBCDIC and write a message to stderr
    * before exiting the program.
    */
    if (JNI_CreateJavaVM("myJVM, (void **)myEnv, (void *)"initArgs)) {
#pragma convert(0)
        fprintf(stderr, "Failed to create the JVM\n");
#pragma convert(819)
        exit(1);
    }

    /* Use the newly created JVM to find the example class,
    * called 'NativeHello'.
    */
    myClass = (*myEnv)->FindClass(myEnv, "NativeHello");
    if (! myClass) {
        Oops(myEnv, "Failed to find class 'NativeHello'");
    }

    /* Now, get the method identifier for the 'main' entry point
    * of the class.
    * Note: The signature of 'main' is always the same for any
    * class called by the following java command:
    * "main" , "([Ljava/lang/String;)V"
    */
    mainID = (*myEnv)->GetStaticMethodID(myEnv,myClass,"main",
        "([Ljava/lang/String;)V");

    if (! mainID) {
        Oops(myEnv, "Failed to find jmethodID of 'main'");
    }
}

```

```

/* Get the jclass for String to create the array
 * of String to pass to 'main'.
 */
stringClass = (*myEnv)->FindClass(myEnv, "java/lang/String");
if (! stringClass) {
    Oops(myEnv, "Failed to find java/lang/String");
}

/* Now, you need to create an empty array of strings,
 * since main requires such an array as a parameter.
 */
args = (*myEnv)->NewObjectArray(myEnv,0,stringClass,0);
if (! args) {
    Oops(myEnv, "Failed to create args array");
}

/* Now, you have the methodID of main and the class, so you can
 * call the main method.
 */
(*myEnv)->CallStaticVoidMethod(myEnv,myClass,mainID,args);

/* Check for errors. */
if ((*myEnv)->ExceptionOccurred(myEnv)) {
    (*myEnv)->ExceptionDescribe(myEnv);
}

/* Finally, destroy the JVM that you created. */
(*myJVM)->DestroyJavaVM(myJVM);

/* All done. */
return 0;
}

```



For more information, see [Java Invocation API](#).

Java native methods and threads considerations

You can use native methods to access functions that are not available in Java^(TM).

To better use Java with native methods, you need to understand these concepts:

- A Java thread, whether created by Java or an attached native thread, has all floating point exceptions disabled. If the thread runs a native method that reenables floating point exceptions, Java does not turn them off a second time. If the user application does not disable them before returning to run Java code, then the Java code may not behave correctly if a floating point exception occurs. When a native thread detaches from the Java virtual machine, its floating point exception mask is restored to the value that was in effect when it was attached.
- When a native thread attaches to the Java virtual machine, the Java virtual machine changes the threads priority, if necessary, to conform to the one to ten priority schemes that Java defines. When the thread detaches, the priority is restored. After attaching, the thread can change the thread priority by using a native method interface (for example, a POSIX API). Java does not change the thread priority on transitions back to the Java virtual machine.
- The Invocation API component of the Java Native Interface (JNI) permits a user to embed a Java virtual machine within their application. If an application creates a Java virtual machine and the Java virtual machine ends abnormally, the MCH74A5 "Java Virtual Machine Terminated" iSeries exception is signalled to the initial thread of the process if that thread was attached to the Java virtual machine when the Java virtual machine ended. The Java virtual machine could end abnormally for any of these reasons:
 - The user calls the `java.lang.System.exit()` method.
 - A thread that the Java virtual machine requires ends.

- An internal error occurs in the Java virtual machine.

This behavior differs from most other Java platforms. On most other platforms, the process that automatically creates the Java virtual machine ends abruptly as soon as the Java virtual machine ends. If the application monitors and handles a signalled MCH74A5 exception, it may continue to run. Otherwise, the process ends when the exception goes unhandled. By adding the code that deals with the iSeries server-specific MCH74A5 exception, you can make the application less portable to other platforms.

Because native methods always run in a multithreaded process, the code that they contain must be thread safe. This places these restrictions on the languages and functions that are used for native methods:

- You should not use ILE CL for native methods, because this language is not thread safe. To run thread safe CL commands, you can use the C language system() function or the java.lang.Runtime.exec() method.
 - Use the C language system() function to run thread safe CL commands from within a C or C++ native method.
 - Use the java.lang.Runtime.exec() method to run thread safe CL commands directly from Java.
- You can use ILE C, ILE C++, ILE COBOL, and ILE RPG to write a native method, but all of the functions that are called from within the native method must be thread safe.

Note: Compile-time support for writing native methods is currently only supplied for the C and C++ languages. While possible, writing native methods in other languages may be much more complicated.

Caution:

Not all standard C, C++, COBOL, or RPG functions are thread safe.

- The C and C++ exit() and abort() functions should never be used within a native method. These functions cause the entire process that runs the Java virtual machine to stop. This includes all of the threads in the process, regardless of if they were originated by Java or not.

Note: The exit() function referred to is the C and C++ function, and is not the same as the java.lang.Runtime.exit() method.

For more information about threads on the iSeries server, see *Multithreaded applications*.

Native methods and the Java Native Interface

Native methods are Java[™] methods that start in a language other than Java. Native methods can access system-specific functions and APIs that are not available directly in Java.

The use of native methods limits the portability of an application, because it involves system-specific code. Native methods can either be new native code statements or native code statements that call existing native code.

Once you decide that a native method is required, it may have to interoperate with the Java virtual machine where it runs. The Java Native Interface (JNI) facilitates this interoperability in a platform-neutral way.

The JNI is a set of interfaces that permit a native method to interoperate with the Java virtual machine in numerous ways. For example, the JNI includes interfaces that create new objects and call methods, get fields and set fields, process exceptions, and manipulate strings and arrays.

For a complete description of the JNI, refer to the Java Native Interface by Sun Microsystems, Inc., or The Source for Java Technology java.sun.com  .

Strings in native methods

Many Java[™] Native Interface (JNI) functions accept C language-style strings as parameters. For example, the FindClass() JNI function accepts a string parameter that specifies the fully-qualified name of a classfile. If the classfile is found, it is loaded by FindClass, and a reference to it is returned to the caller of FindClass.

All JNI functions expect their string parameters to be encoded in UTF-8. For details on UTF-8, you can refer to the JNI Specification, but in most cases it is enough to observe that 7-bit American Standard Code for Information Interchange (ASCII) characters are equivalent to their UTF-8 representation. 7-bit ASCII characters are actually 8-bit characters but their first bit is always 0. So, most ASCII C strings are actually already in UTF-8.

The C compiler on the iSeries server operates in extended binary-coded decimal interchange code (EBCDIC) by default, so you can provide strings to the JNI functions in UTF-8. There are two ways to do this. You can use literal strings, or you can use dynamic strings. Literal strings are strings whose value is known when the source code is compiled. Dynamic strings are strings whose value is not known at compile time, but is actually computed at run time.

Literal strings in native methods: It is easier to encode literal strings in UTF-8 if the string is composed of characters with a 7-bit American Standard Code for Information Interchange (ASCII) representation. If the string can be represented in ASCII, as most are, then the string can be bracketed by 'pragma' statements that change the current codepage of the compiler. Then, the compiler stores the string internally in the UTF-8 form that is required by the JNI. If the string cannot be represented in ASCII, it is easier to treat the original extended binary-coded decimal interchange code (EBCDIC) string as a dynamic string, and process it using `iconv()` before passing it to the JNI. For more information on dynamic strings, see dynamic strings.

For example, to find the class named `java/lang/String`, the code looks like this:

```
#pragma convert(819)
myClass = (*env)->FindClass(env,"java/lang/String");
#pragma convert(0)
```

The first pragma, with the number 819, informs the compiler to store all subsequent double-quoted strings (literal strings) in ASCII. The second pragma, with the number 0, tells the compiler to revert to the default code page of the compiler for double-quoted strings, which is usually the EBCDIC code page 37. So, by bracketing this call with these pragmas, we satisfy the JNI requirement that string parameters are encoded in UTF-8.

Caution: Be careful with text substitutions. For example, if your code looks like this:

```
#pragma convert(819)
#define MyString "java/lang/String"
#pragma convert(0)
myClass = (*env)->FindClass(env,MyString);
```

Then, the resulting string is EBCDIC, because the value of `MyString` is substituted into the `FindClass` call during compilation. At the time of this substitution, the pragma, number 819, is not in effect. Thus, literal strings are not stored in ASCII.

Convert dynamic strings to and from EBCDIC, Unicode, and UTF-8: To manipulate string variables that are computed at run time, it may be necessary to convert strings to and from extended binary-coded decimal interchange (EBCDIC), Unicode, and UTF-8.

The system API that provides for code page conversion function is `iconv()`. To use `iconv()`, follow these steps:

1. Create a conversion descriptor with `QtqIconvOpen()`.
2. Call `iconv()` to use the descriptor to convert to a string.
3. Close the descriptor by using `iconv_close`.

In Example 3 of the using the Java^(TM) Native Interface for native methods examples, the routine creates, uses, and then destroys the `iconv` conversion descriptor within the routine. This scheme avoids the

problems with multithreaded use of an `iconv_t` descriptor, but for performance sensitive code it is better to create a conversion descriptor in static storage, and moderate multiple access to it using a mutual exclusion (mutex) or other synchronization facility.

Examples: Use the Java Native Interface for native methods

This example program is a simple Java[™] Native Interface (JNI) example in which a C native method is used to display "Hello, World." Use the `javah` tool with the `NativeHello` class file to generate the `NativeHello.h` file. This example assumes that the `NativeHello` C implementation is part of a service program that is called `NATHELLO`.

Note: The library where the `NATHELLO` service program is located must be in the library list for this example to run.

Example 1: NativeHello class

Note: Read the Code example disclaimer for important legal information.

```
public class NativeHello {

    // Declare a field of type 'String' in the NativeHello object.
    // This is an 'instance' field, so every NativeHello object
    // contains one.
    public String theString;          // instance variable

    // Declare the native method itself. This native method
    // creates a new string object, and places a reference to it
    // into 'theString'
    public native void setTheString(); // native method to set string

    // This 'static initializer' code is called before the class is
    // first used.
    static {

        // Attempt to load the native method library. If you do not
        // find it, write a message to 'out', and try a hardcoded path.
        // If that fails, then exit.
        try {

            // System.loadLibrary uses the iSeries library list in JDK 1.1,
            // and uses the java.library.path property or the LIBPATH environment
            // variable in JDK1.2
            System.loadLibrary("NATHELLO");
        }

        catch (UnsatisfiedLinkError e1) {

            // Did not find the service program.
            System.out.println
                ("I did not find NATHELLO *SRVPGM.");
            System.out.println ("(I will try a hardcoded path)");

            try {

                // System.load takes the full integrated file system form path.
                System.load ("/qsys.lib/jniexample.lib/nathello.srvpgm");
            }

            catch (UnsatisfiedLinkError e2) {

                // If you get to this point, then you are done! Write the message
                // and exit.
                System.out.println
                    ("<sigh> I did not find NATHELLO *SRVPGM anywhere. Goodbye");
                System.exit(1);
            }
        }
    }
}
```

```

    }
}

// Here is the 'main' code of this class. This is what runs when you
// enter 'java NativeHello' on the command line.
public static void main(String argv[]){

    // Allocate a new NativeHello object now.
    NativeHello nh = new NativeHello();

    // Echo location.
    System.out.println("(Java) Instantiated NativeHello object");
    System.out.println("(Java) string field is '" + nh.theString + "'");
    System.out.println("(Java) Calling native method to set the string");

    // Here is the call to the native method.
    nh.setTheString();

    // Now, print the value after the call to double check.
    System.out.println("(Java) Returned from the native method");
    System.out.println("(Java) string field is '" + nh.theString + "'");
    System.out.println("(Java) All done...");
}
}

```

Example 2: Generated NativeHello.h header file

Note: Read the Code example disclaimer for important legal information.

```

/* DO NOT EDIT THIS FILE - it is machine generated */
#include <jni.h>
/* Header for class NativeHello */

#ifdef _Included_NativeHello
#define _Included_NativeHello
#ifdef __cplusplus
extern "C" {
#endif
/*
 * Class:      NativeHello
 * Method:     setTheString
 * Signature:  ()V
 */
JNIEXPORT void JNICALL Java_NativeHello_setTheString
    (JNIEnv *, jobject);

#ifdef __cplusplus
}
#endif
#endif

```

This NativeHello.c example shows the implementation of the native method in C. This example shows how to link Java to native methods. However, it points out complications that arise from the fact that the iSeries server is internally an extended binary-coded decimal interchange code (EBCDIC) machine. It also shows complications from the current lack of true internationalization elements in the JNI.

These reasons, although they are not new with the JNI, cause some unique iSeries server-specific differences in the C code that you write. You must remember that if you are writing to stdout or stderr or reading from stdin, your data is probably encoded in EBCDIC form.

In C code, you can easily convert most literal strings, those that contain 7-bit characters only, into the UTF-8 form that is required by the JNI. To do this, bracket the literal strings with code-page conversion pragmas. However, because you may write information directly to stdout or stderr from your C code, you might allow some literals to remain in EBCDIC.

Note: The `#pragma convert(0)` statements convert character data to EBCDIC. The `#pragma convert(819)` statements convert character data to American Standard Code for Information Interchange (ASCII). These statements convert character data in the C program at compile time.

Example 3: NativeHello.c native method implementation of the NativeHello Java class

Note: Read the Code example disclaimer for important legal information.

```
#include <stdlib.h>      /* malloc, free, and so forth */
#include <stdio.h>      /* fprintf(), and so forth */
#include <qtqiconv.H>   /* iconv() interface */
#include <string.h>     /* memset(), and so forth */
#include "NativeHello.h" /* generated by 'javah-jni' */

/* All literal strings are ISO-8859-1 Latin 1 code page (and with 7-bit
characters, they are also automatically UTF-8). */
#pragma convert(819) /* handle all literal strings as ASCII */

/* Report and clear a JNI exception. */
static void HandleError(JNIEnv*);

/* Print an UTF-8 string to stderr in the coded character */
/* set identifier (CCSID) of the current job. */
static void JobPrint(JNIEnv*, char*);

/* Constants describing which direction to convert: */
#define CONV_UTF2JOB 1
#define CONV_JOB2UTF 2

/* Convert a string from the CCSID of the job to UTF-8, or vice-versa. */
int StringConvert(int direction, char *sourceStr, char *targetStr);

/* Native method implementation of 'setTheString()'. */

JNIEXPORT void JNICALL Java_NativeHello_setTheString
(JNIEnv *env, jobject javaThis)
{
    jclass thisClass; /* class for 'this' object */
    jstring stringObject; /* new string, to be put in field in 'this' */
    jfieldID fid; /* field ID required to update field in 'this' */
    jthrowable exception; /* exception, retrieved using ExceptionOccurred */

    /* Write status to console. */
    JobPrint(env, "( C ) In the native method\n");

    /* Build the new string object. */
    if (! (stringObject = (*env)->NewStringUTF(env, "Hello, native world!")))
    {
        /* For nearly every function in the JNI, a null return value indicates
        that there was an error, and that an exception had been placed where it
        could be retrieved by 'ExceptionOccurred()'. In this case, the error
        would typically be fatal, but for purposes of this example, go ahead
        and catch the error, and continue. */
        HandleError(env);
        return;
    }

    /* get the class of the 'this' object, required to get the fieldID */
    if (! (thisClass = (*env)->GetObjectClass(env, javaThis)))
    {
        /* A null class returned from GetObjectClass indicates that there
        was a problem. Instead of handling this problem, simply return and
        know that the return to Java automatically 'throws' the stored Java
        exception. */
        return;
    }
}
```

```

/* Get the fieldID to update. */
if (! (fid = (*env)->GetFieldID(env,
                               thisClass,
                               "theString",
                               "Ljava/lang/String;")))
{
    /* A null fieldID returned from GetFieldID indicates that there
    was a problem. Report the problem from here and clear it.
    Leave the string unchanged. */
    HandleError(env);
    return;
}

JobPrint(env, "( C ) Setting the field\n");

/* Make the actual update.
Note: SetObjectField is an example of an interface that does
not return a return value that can be tested. In this case, it
is necessary to call ExceptionOccurred() to see if there
was a problem with storing the value */
(*env)->SetObjectField(env, javaThis, fid, stringObject);

/* Check to see if the update was successful. If not, report the error. */
if ((*env)->ExceptionOccurred(env)) {
    /* A non-null exception object came back from ExceptionOccurred,
    so there is a problem and you must report the error. */
    HandleError(env);
}

JobPrint(env, "( C ) Returning from the native method\n");
return;
}

static void HandleError(JNIEnv *env)
{
    /* A simple routine to report and handle an exception. */
    JobPrint(env, "( C ) Error occurred on JNI call: ");
    (*env)->ExceptionDescribe(env); /* write exception data to the console */
    (*env)->ExceptionClear(env); /* clear the exception that was pending */
}

static void JobPrint(JNIEnv *env, char *str)
{
    char *jobStr;
    char buf[512];
    size_t len;

    len = strlen(str);

    /* Only print non-empty string. */
    if (len) {
        jobStr = (len >= 512) ? malloc(len+1) : &buf;
        if (! StringConvert(CONV_UTF2JOB, str, jobStr))
            (*env)->FatalError
            (env, "ERROR in JobPrint: Unable to convert UTF2JOB");
        fprintf(stderr, jobStr);
        if (len >= 512) free(jobStr);
    }
}

int StringConvert(int direction, char *sourceStr, char *targetStr)
{
    QtqCode_T source, target; /* parameters to instantiate iconv */
    size_t sStrLen, tStrLen; /* local copies of string lengths */

```

```

iconv_t  ourConverter;      /* the actual conversion descriptor */
int      iconvRC;          /* return code from the conversion */
size_t   originalLen;      /* original length of the sourceStr */

/* Make local copies of the input and output sizes that are initialized
to the size of the input string. The iconv() requires the
length parameters to be passed by address (that is as int*). */
originalLen = sStrLen = tStrLen = strlen(sourceStr);

/* Initialize the parameters to the QtqIconvOpen() to zero. */
memset(&source,0x00,sizeof(source));
memset(&target,0x00,sizeof(target));

/* Depending on direction parameter, set either SOURCE
or TARGET CCSID to ISO 8859-1 Latin. */
if (CONV_UTF2JOB == direction ) {
    source.CCSID = 819;
}
else {
    target.CCSID = 819;
}

/* Create the iconv_t converter object. */
ourConverter = QtqIconvOpen(&target,&source);

/* Make sure that you have a valid converter, otherwise return 0. */
if (-1 == ourConverter.return_value) return 0;

/* Perform the conversion. */
iconvRC = iconv(ourConverter,
               (char**) &sourceStr,
               &sStrLen,
               &targetStr,
               &tStrLen);

/* If the conversion failed, return a zero. */
if (0 != iconvRC ) return 0;

/* Close the conversion descriptor. */
iconv_close(ourConverter);

/* The targetStr returns pointing to the character just
past the last converted character, so set the null
there now. */
*targetStr = '\0';

/* Return the number of characters that were processed. */
return originalLen-tStrLen;
}

#pragma convert(0)

```

See [Use the Java Native Interface for native methods](#) for background information. 

IBM OS/400 PASE native methods for Java

The iSeries Java^(TM) virtual machine (JVM) now supports the use of native methods running in the OS/400^(R) PASE environment. Prior to V5R2, the native iSeries JVM used only ILE native methods. Support for OS/400 PASE native methods includes:

- Full use of the native iSeries Java Native Interface (JNI) from OS/400 PASE native methods
- The ability to call OS/400 PASE native methods from the native iSeries JVM

This new support enables you to easily port your Java applications that run in AIX^(R) to your iSeries server. You can copy the class files and AIX native method libraries to the integrated file system on the iSeries and run them from any of the control language (CL), Qshell or OS/400 PASE terminal session command prompts.

For more information about using IBM OS/400 PASE native methods for Java, see the following topics:

Java OS/400 PASE environment variables

Learn about the environment variables that you must define before using OS/400 PASE native methods. These environment variables manage the OS/400 PASE and JVM runtime environments.

Java OS/400 PASE error codes

To help you troubleshoot OS/400 PASE native methods, find out about error conditions that are indicated by OS/400 job log messages and Java runtime exceptions.

Managing native method libraries

Find out about the Java library naming conventions and the library search algorithm. This information is important for managing multiple versions of a native method library on your iSeries server.

Example: IBM OS/400 PASE native methods for Java

See how to run a simple Java program that prints out the contents of a Java String. Rather than accessing the string directly from Java code, the example calls a native method that then calls back into Java through JNI to get the string value.

This information assumes you are already familiar with OS/400 PASE. For more information, see the following topic:

OS/400 PASE



Java OS/400 PASE environment variables

The Java virtual machine (JVM) uses the following variables to start OS/400 PASE environments. You need to set the QIBM_JAVA_PASE_STARTUP variable in order to run the IBM OS/400 PASE native method for Java example.

For information about setting environment variables for the example, see the following topic:

Environment variables for the IBM OS/400 PASE example.

QIBM_JAVA_PASE_STARTUP

You need to set this environment variable when both of the following conditions occur:

- You are using OS/400 PASE native methods
- You are starting Java from an iSeries command prompt or Qshell command prompt

The JVM uses this environment variable to start a PASE environment. The value of the variable identifies an OS/400 PASE startup program. Your iSeries server includes two OS/400 PASE startup programs:

- /usr/lib/start32: Starts a 32-bit OS/400 PASE environment
- /usr/lib/start64: Starts a 64-bit OS/400 PASE environment

The bit format of all shared library objects used by an OS/400 PASE environment must match the bit format of the OS/400 PASE environment.

You can not use this variable when starting Java from an OS/400 PASE terminal session. An OS/400 PASE terminal session always uses a 32-bit OS/400 PASE environment. Any JVMs started from an OS/400 PASE terminal session use the same type of PASE environment as the terminal session.

QIBM_JAVA_PASE_CHILD_STARTUP

Set this optional environment variable when the OS/400 PASE environment for a secondary JVM must be different than the OS/400 PASE environment the primary JVM. A call to `Runtime.exec()` in Java starts a secondary (or child) JVM.

For more information, see [Using QIBM_JAVA_PASE_CHILD_STARTUP](#).



Examples: Environment variables for the IBM OS/400 PASE example: To use the IBM OS/400 PASE native methods for Java example, you need to set the following environment variables.

PASE_LIBPATH

Your iSeries server uses this OS/400 PASE environment variable to identify the location of OS/400 PASE native method libraries. You can set the path to a single directory or to multiple directories. For multiple directories, use a colon (:) to separate entries. Your server can also use the LIBPATH environment variable.

For more information about using Java, native method libraries, and PASE_LIBPATH with this example, see the following topic:

[Using Java, OS/400 PASE, and native method libraries](#)

PASE_THREAD_ATTACH

Setting this OS/400 PASE environment variable to Y causes an ILE thread that was not started by OS/400 PASE to be attached automatically to OS/400 PASE when it calls an OS/400 PASE procedure.

For more information about OS/400 PASE environment variables, see the appropriate entries in the following topic:

[Work with OS/400 PASE environment variables](#)

QIBM_JAVA_PASE_STARTUP

The JVM uses this environment variable to start an OS/400 PASE environment. The value of the variable identifies an OS/400 PASE startup program.

For more information, see the following topic:

[Java OS/400 PASE variables](#)



Using QIBM_JAVA_PASE_CHILD_STARTUP: The QIBM_JAVA_PASE_CHILD_STARTUP environment variable indicates the OS/400 PASE startup program for any secondary JVMs. Use QIBM_JAVA_PASE_CHILD_STARTUP when all of the following conditions are true:

- The Java application that you want to run creates Java virtual machines (JVMs) through Java calls to `Runtime.exec()`
- Both the primary and secondary JVMs use OS/400 PASE native methods
- The OS/400 PASE environment of the secondary JVMs must be different than the OS/400 PASE environment of the primary JVM

When all of the previously listed conditions are true, perform the following actions:

- Set the QIBM_JAVA_PASE_CHILD_STARTUP environment variable to the OS/400 PASE startup program of the secondary JVMs
- When starting the primary JVM from an iSeries command prompt or Qshell command prompt, set the QIBM_JAVA_PASE_STARTUP environment variable to the OS/400 PASE startup program of the primary JVM.

Note: When starting the primary JVM from an OS/400 PASE terminal session, do not set QIBM_JAVA_PASE_STARTUP.

The process of the secondary JVM inherits the QIBM_JAVA_PASE_CHILD_STARTUP environment variable. In addition, OS/400 sets the QIBM_JAVA_PASE_STARTUP environment variable of the secondary JVM process to the value of the QIBM_JAVA_PASE_CHILD_STARTUP environment variable from the parent process.

The following table identifies the resulting OS/400 PASE environments (if any) for the various combinations of command environments and definitions of QIBM_JAVA_PASE_STARTUP and QIBM_JAVA_PASE_CHILD_STARTUP:

Starting Environment			Resulting Behavior	
Command environment	QIBM_JAVA_PASE_STARTUP	QIBM_JAVA_PASE_CHILD_STARTUP	Primary JVM OS/400 PASE Startup	Secondary JVM OS/400 PASE Startup
CL or QSH	Defined startX	Defined startY	Use startX	Use startY
CL or QSH	Defined startX	Not defined	Use startX	Use startX
CL or QSH	Not defined	Defined startY	No OS/400 PASE environment	Use startY
CL or QSH	Not defined	Not defined	No OS/400 PASE environment	No OS/400 PASE environment
OS/400 PASE terminal session	Defined startX	Defined startY	Not Allowed*	Not Allowed*
OS/400 PASE terminal session	Defined startX	Not defined	Not Allowed*	Not Allowed*
OS/400 PASE terminal session	Not defined	Defined startY	Use OS/400 PASE terminal session environment	Use startY
OS/400 PASE terminal session	Not defined	Not defined	Use OS/400 PASE terminal session environment	No OS/400 PASE environment

* The rows marked 'Not allowed' indicate situations where the QIBM_JAVA_PASE_STARTUP environment variable could conflict with the OS/400 PASE terminal session. Because of the potential conflict, using QIBM_JAVA_PASE_STARTUP is not allowed from an OS/400 PASE terminal session. >>

Managing native method libraries

To use native method libraries, especially when you want to manage multiple versions of a native method library on your iSeries server, you need to understand both the Java library naming conventions and the library search algorithm.

OS/400 uses the first native method library that matches the name of the library that the Java virtual machine (JVM) loads. In order to ensure that OS/400 finds the correct native methods, you must avoid library name clashes and confusion about which native method library the JVM uses.

OS/400 PASE and AIX Java Library Naming Conventions: If the Java code loads a library named Sample, the corresponding executable file must be named either libSample.a or libSample.so.

Java library search order: When you enable OS/400 PASE native methods for the JVM, your server uses three different lists (in the following order) to create a single native method library search path:

1. OS/400 library list
2. LIBPATH environment variable
3. PASE_LIBPATH environment variable

In order to perform the search, OS/400 converts the library list to the integrated file system format. QSYS file system objects have equivalent names in the integrated file system, but some integrated file system objects do not have equivalent QSYS file system names. Because the library loader looks for objects in both the QSYS file system and in the integrated file system, OS/400 uses the integrated file system format to search for native method libraries.

The following table shows how OS/400 converts entries in the library list to the integrated file system format:

Library list entry	Integrated file system format
QSYS	/qsys.lib
QSYS2	/qsys.lib/qsys2.lib
QGPL	/qsys.lib/qgpl.lib
QTEMP	/qsys.lib/qtemp.lib

Example: Searching for the Sample2 library

In the following example, LIBPATH is set to /home/user1/lib32:/samples/lib32 and PASE_LIBPATH is set to /QOpenSys/samples/lib.

The following table, when read from top to bottom, indicates the full search path:

Source	Integrated file system directories
Library list	/qsys.lib /qsys.lib/qsys2.lib /qsys.lib/qgpl.lib /qsys.lib/qtemp.lib
LIBPATH	/home/user1/lib32 /samples/lib32
PASE_LIBPATH	/QOpenSys/samples/lib

Note: Uppercase and lowercase characters are significant only in the /QOpenSys path.

In order to search for library Sample2, the Java library loader searches for file candidates in the following order:

1. /qsys.lib/sample2.srvpgm
2. /qsys.lib/libSample2.a
3. /qsys.lib/libSample2.so
1. /qsys.lib/qsys2.lib/sample2.srvpgm
2. /qsys.lib/qsys2.lib/libSample2.a
3. /qsys.lib/qsys2.lib/libSample2.so

1. /qsys.lib/qgpl.lib/sample2.srvpgm
2. /qsys.lib/qgpl.lib/libSample2.a
3. /qsys.lib/qgpl.lib/libSample2.so
1. /qsys.lib/qtemp.lib/sample2.srvpgm
2. /qsys.lib/qtemp.lib/libSample2.a
3. /qsys.lib/qtemp.lib/libSample2.so
1. /home/user1/lib32/sample2.srvpgm
2. /home/user1/lib32/libSample2.a
3. /home/user1/lib32/libSample2.so
1. /samples/lib32/sample2.srvpgm
2. /samples/lib32/libSample2.a
3. /samples/lib32/libSample2.so
1. /QOpenSys/samples/lib/SAMPLE2.srvpgm
2. /QOpenSys/samples/lib/libSample2.a
3. /QOpenSys/samples/lib/libSample2.so

OS/400 loads the first candidate in the list that actually exists into the JVM as a native method library. Even though candidates like 'qsys.lib/libSample2.a' and 'qsys.lib/libSample2.so' occur in the search, it is not possible to create integrated file system files or symbolic links in the /qsys.lib directories. Therefore, even though OS/400 checks for these candidate files, it will never find them in integrated file system directories that begin with /qsys.lib.

However, you can create arbitrary symbolic links from other integrated file system directories to OS/400 objects in the QSYS file system. As a result, valid [»](#)

Java OS/400 PASE error codes

The following lists describe errors that you may encounter at start up or at run time when using OS/400 PASE native methods for Java.

Startup Errors: There are 3 new error codes for the JVAB55C “Unable to create Java virtual machine” message:

- 19 - Error starting OS/400 PASE environment. Indicates a user application or operating system problem. Error code 19 also includes English-only text. You may see the following error text:
 - Java OS/400 PASE error. OS/400 PASE is already active and the QIBM_JAVA_PASE_STARTUP environment variable is defined.

Either remove the QIBM_JAVA_PASE_STARTUP environment variable definition or end the active OS/400 PASE terminal session.
 - Java OS/400 PASE error. Unable to run OS/400 PASE startup program &programName.

Either the OS/400 PASE program identified by the QIBM_JAVA_PASE_STARTUP environment variable does not exist, or you can not run the program in the OS/400 PASE environment.
 - Java OS/400 PASE internal error number &errorCode.

You may see any of the following internal error numbers:

 - 106 - OS/400 PASE or the specified OS/400 PASE format is not supported for the indicated JDK version.

Either OS/400 PASE is not supported for the specified JDK, or else the OS/400 PASE startup program bit format is not supported for the specified JDK. For V5R2 the supported combinations are:

- JDK 1.2 and OS/400 PASE 32-bit format
- JDK 1.3 and OS/400 PASE 32-bit format
- JDK 1.3 and OS/400 PASE 64-bit format

Report the following error codes to a service representative:

- 101 - No startup program identified.
- 102 - Unable to retrieve OS/400 PASE JavaVM pointer
- 103 - Unable to locate Qp2CallPase
- 104 - OS/400 PASE pointer size error.
- 105 - Cannot locate OS/400 PASE libjvm.a.
- 20 - OS/400 PASE operand not valid. Report to a service representative.
- 21 - Unable to attach job to OS/400 PASE. Report to a service representative.

Runtime errors: In addition to startup errors, PASEInternalError or PASEExit Java exceptions may appear in the Qshell output of the JVM:

- PASEInternalError - indicates internal system error. Check for Licensed Internal Code Log entries. For more information, see Qp2CallPase.
- PASEExit - either the OS/400 PASE application called the exit() function or the OS/400 PASE environment ended abnormally. Check the Job Log and Licensed Internal Code Log for additional information. <<



Example: IBM OS/400 PASE native method for Java

The IBM OS/400 PASE native method for Java example calls an instance of a native C method that then uses Java Native Interface (JNI) to call back into Java code.

To see HTML versions of the example source files, use the following links:

- PASEExample1.java
- PASEExample1.c

Before you can run the OS/400 PASE native method example, you must complete the following tasks:

1. Download the example source code to your AIX workstation
2. Prepare the example source code
3. Prepare your iSeries server

Run the OS/400 PASE native method for Java example: After you complete the previous tasks, you can run the example. Use either of the following commands to run the example program:

- From an iSeries server command prompt:

```
JAVA CLASS(PASEExample1) CLASSPATH('/home/example')
```

- From a Qshell command prompt or OS/400 PASE terminal session:

```
cd /home/example
java PASEExample1
```

Comparison of Integrated Language Environment and Java

The Java[™] environment on an iSeries server is separate from the integrated language environment (ILE). Java is not an ILE language, and it cannot bind to ILE object modules to create programs or service programs on an iSeries server.

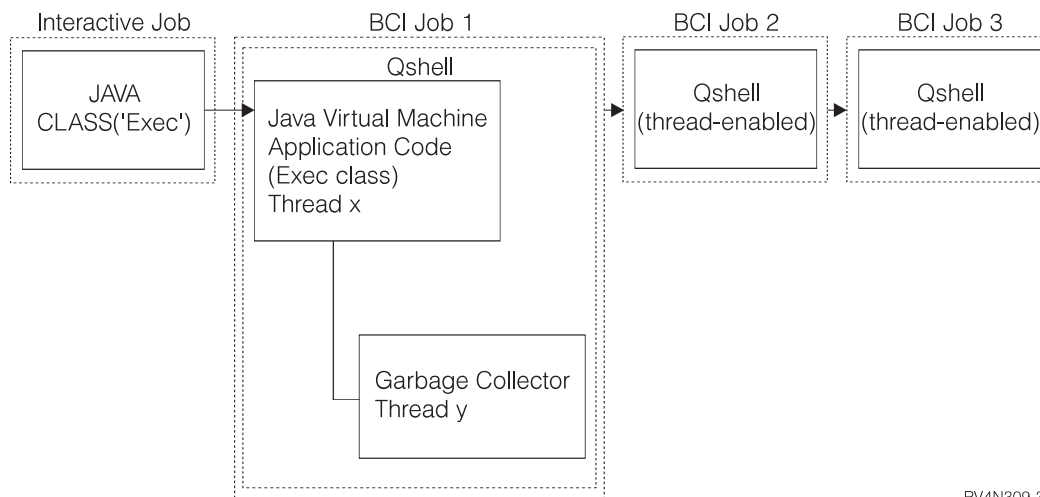
ILE	Java
Members that are part of the library or file structure on an iSeries server store source codes.	Stream files in the integrated file system contain source code.
Source entry utility (SEU) edits extended binary-coded decimal interchange code (EBCDIC) source files.	American Standard Code for Information Interchange (ASCII) source files are usually edited using a workstation editor.
Source files compile into object code modules, which are stored in libraries on an iSeries server.	Source code compiles into class files, which the integrated file system stores.
Object modules are statically bound together in programs or service programs.	Classes are dynamically loaded, as needed, at runtime.
You can directly call to functions that are written in other ILE programming languages.	Java Native Interface must be used to call other languages from Java.
ILE languages are always compiled and run as machine instructions.	Java programs can be interpreted or compiled.

Use `java.lang.Runtime.exec()`

The `java.lang.Runtime.exec()` method calls programs or commands from within a Java[™] program. The actual processing that occurs depends on exactly what information is passed to the `exec()` method. In all cases, the `Runtime.exec()` method creates another thread-enabled batch immediate (BCI) job. The BCI job processes the command string that is passed in on the `Runtime.exec()` method.

To use the `java.lang.Runtime.exec()` method, you must install the Qshell Interpreter on your iSeries server if you are using JDK 1.2 or lower. For more information about the Qshell Interpreter, see Qshell Interpreter.

Note: The `java.lang.Runtime.exec()` method runs programs in a separate process. This differs from the C system function, which runs a program in the same process.



RV4N309-3

If the command being processed is a Qshell utility, it runs in the second BCI job, and the third BCI job is not created. If the command being processed is a CL command, the second BCI job is started to run the Qshell, and the third BCI job is started to run the CL command. A Qshell utility is a built-in utility that QSH can run directly. An example of a Qshell utility is the `javac` command, which compiles Java programs. The processing in the second (or third) BCI job runs concurrently with the Java virtual machine. Any exit or shutdown processing in those jobs does not affect the original Java virtual machine.

When calling an iSeries command or program, you must ensure that any parameters being passed to the called program are in the code page that is expected by that program.

See call another Java program, call a CL program, or call a CL command for `java.lang.Runtime.exec()` examples.

Example: Call another Java program with `java.lang.Runtime.exec()`

This example shows how to call another Java^(TM) program with `java.lang.Runtime.exec()`. This class calls the Hello program that is shipped as part of the IBM Developer Kit for Java. When the Hello class writes to `System.out`, this program gets a handle to the stream and can read from it.

Note: You use the Qshell Interpreter to call the program.

Example 1: CallHelloPgm class

Note: Read the Code example disclaimer for important legal information.

```
import java.io.*;

public class CallHelloPgm
{
    public static void main(String args[])
    {
        Process theProcess = null;
        BufferedReader inStream = null;

        System.out.println("CallHelloPgm.main() invoked");

        // call the Hello class
        try
        {
            theProcess = Runtime.getRuntime().exec("java com.ibm.as400.system.Hello");
        }
        catch(IOException e)
        {
            System.err.println("Error on exec() method");
            e.printStackTrace();
        }

        // read from the called program's standard output stream
        try
        {
            inStream = new BufferedReader(
                new InputStreamReader( theProcess.getInputStream() ));
            System.out.println(inStream.readLine());
        }
        catch(IOException e)
        {
            System.err.println("Error on inStream.readLine()");
            e.printStackTrace();
        }

    } // end method
} // end class
```

For background information, see `Use java.lang.Runtime.exec()`.

Example: Call a CL program with `java.lang.Runtime.exec()`

This example shows how to run CL programs from within a Java^(TM) program. See call a CL command for an example of how to call a CL command from within a Java program. In this example, the Java class `CallCLPgm` runs a CL program. The CL program uses the `Display Java Program (DSPJVAPGM)` command to display the program that is associated with the Hello class file. This example assumes that the CL program has been compiled and exists in a library that is called `JAVSAMPLIB`. The output from the CL program is in the `QSYSPRT` spooled file.

Note: The JAVSAMPLIB is not created as part of the IBM Developer Kit licensed program (LP) number 5722-JV1 installation process. You must explicitly create the library.

Example 1: CallCLPgm class

Note: Read the Code example disclaimer for important legal information.

```
import java.io.*;

public class CallCLPgm
{
    public static void main(String[] args)
    {
        try
        {
            Process theProcess =
                Runtime.getRuntime().exec("/QSYS.LIB/JAVSAMPLIB.LIB/DSPJVA.PGM");
        }
        catch(IOException e)
        {
            System.err.println("Error on exec() method");
            e.printStackTrace();
        }
    } // end main() method
} // end class
```

Example 2: Display Java CL program

```
PGM
DSPJVAPGM CLSF('/QIBM/ProdData/Java400/com/ibm/as400/system/Hello.class') +
          OUTPUT(*PRINT)
ENDPGM
```

For background information, see Use java.lang.Runtime.exec().

Example: Call a CL command with java.lang.Runtime.exec()

This example shows how to run a control language (CL) command from within a Java program. In this example, the Java class runs a CL command. The CL command uses the Display Java Program (DSPJVAPGM) command to display the program that is associated with the Hello class file. The output from the CL command is in the QSYSPRT spooled file.

» If you are using JDK 1.1.8 or JDK 1.2, « every command that is passed into the Runtime.getRuntime().exec() function must be in quotation marks and in Qshell format. In addition, to run a CL command from Qshell, it is necessary to pass in the string:

```
"system \"CL COMMAND\""
```

Where *CL COMMAND* is the CL command you want to run. So the line to call the command MYCLCOM should be this:

```
Runtime.getRuntime().Exec("system \"MYCLCOM\"");
```

» **Note:** When you are using JDK 1.3 or JDK 1.4, omit the slash and quote delimiters (\"). For example, when using JDK version 1.3 or higher, the call to command MYCLCOM looks like this:

```
Runtime.getRuntime().Exec("system MYCLCOM");
```

For more information, see the os400.runtime.exec system property in Java system properties for Java 2 Software Development Kit, Standard Edition. «

Example 1: CallCLCom class

» The following example uses the Qshell delimiters necessary when you are using JDK 1.1.8 or JDK 1.2. If you are using JDK version 1.3 or higher, omit the delimiters. «

Note: Read the Code example disclaimer for important legal information.

```
import java.io.*;

public class CallCLCom
{
    public static void main(String[] args)
    {
        try
        {
            Process theProcess = Runtime.getRuntime().exec("system \"DSPJVAPGM
                CLSF('/com/ibm/as400/system/Hello.class') OUTPUT(*PRINT)\");
        }
        catch(IOException e)
        {
            System.err.println("Error on exec() method");
            e.printStackTrace();
        }
    } // end main() method
} // end class
```

For background information, see `Use java.lang.Runtime.exec()`.

Interprocess communications

When communicating with programs that are running in another process, there are a number of options.

One option is to use sockets for interprocess communication. One program can act as the server program that listens on a socket connection for input from the client program. The client program connects to the server with a socket. Once the socket connection is established, either program can send or receive information.

Another option is to use stream files for communication between programs. To do this, use the `System.in`, `System.out`, and `System.err` classes.

A third option is to use the IBM Toolbox for Java^(TM) which provides data queues and iSeries message objects.

» You can also call Java from other languages. See [Example: Call Java from C](#) and [Example: Call Java from RPG](#) for more information. «

Use sockets for interprocess communication

Sockets streams communicate between programs that are running in separate processes. The programs can either start separately or start by using the `java.lang.Runtime.exec()` method from within the main Java^(TM) program. If a program is written in a language other than Java, you must ensure that any American Standard Code for Information Interchange (ASCII) or extended binary-coded decimal interchange code (EBCDIC) conversion takes place. See [Java character encodings](#) for more details.

For an example that uses sockets, see [Example: Use sockets for interprocess communication](#).

Example: Use sockets for interprocess communication: This example uses sockets to communicate between a Java^(TM) program and a C program. You should start the C program first, which listens on a socket. Once the Java program connects to the socket, the C program sends it a string by using that socket connection. The string that is sent from the C program is an American Standard Code for Information Interchange (ASCII) string in codepage 819.

The Java program should be started using this command, `java TalkToC xxxxx nnnn` on the Qshell Interpreter command line or on another Java platform. Or, enter `JAVA TALKTOC PARM(yyyyy nnnn)` on the iSeries command line to start the Java program. `yyyyy` is the domain name or Internet Protocol (IP) address of the system on which the C program is running. `nnnn` is the port number of the socket that the C program is using. You should also use this port number as the first parameter on the call to the C program.

Example 1: TalkToC client class

Note: Read the Code example disclaimer for important legal information.

```
import java.net.*;
import java.io.*;

class TalkToC
{
    private String host = null;
    private int port = -999;
    private Socket socket = null;
    private BufferedReader inStream = null;

    public static void main(String[] args)
    {
        TalkToC caller = new TalkToC();
        caller.host = args[0];
        caller.port = new Integer(args[1]).intValue();
        caller.setUp();
        caller.converse();
        caller.cleanup();
    } // end main() method

    public void setUp()
    {
        System.out.println("TalkToC.setUp() invoked");

        try
        {
            socket = new Socket(host, port);
            inStream = new BufferedReader(new InputStreamReader(
                socket.getInputStream()));
        }
        catch(UnknownHostException e)
        {
            System.err.println("Cannot find host called: " + host);
            e.printStackTrace();
            System.exit(-1);
        }
        catch(IOException e)
        {
            System.err.println("Could not establish connection for " + host);
            e.printStackTrace();
            System.exit(-1);
        }
    } // end setUp() method

    public void converse()
    {
        System.out.println("TalkToC.converse() invoked");

        if (socket != null && inStream != null)
        {
            try
            {
                System.out.println(inStream.readLine());
            }
        }
    }
}
```



```

        }
        catch(IOException e)
        {
            System.err.println("Conversation error with host " + host);
            e.printStackTrace();
        }

    } // end if

} // end converse() method

public void cleanUp()
{
    try
    {
        if(inStream != null)
        {
            inStream.close();
        }
        if(socket != null)
        {
            socket.close();
        }
    } // end try
    catch(IOException e)
    {
        System.err.println("Error in cleanup");
        e.printStackTrace();
        System.exit(-1);
    }
} // end cleanUp() method

} // end TalkToC class

```

SockServ.C starts by passing in a parameter for the port number. For example, CALL SockServ '2001'.

Example 2: SockServ.C server program

Note: Read the Code example disclaimer for important legal information.

```

#include <stdlib.h>
#include <stdio.h>
#include <errno.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netinet/tcp.h>
#include <unistd.h>
#include <sys/time.h>

void main(int argc, char* argv[])
{
    int    portNum = atoi(argv[1]);
    int    server;
    int    client;
    int    address_len;
    int    sendrc;
    int    bndrc;
    char*  greeting;
    struct sockaddr_in local_Address;
    address_len = sizeof(local_Address);

    memset(&local_Address,0x00,sizeof(local_Address));
    local_Address.sin_family = AF_INET;
    local_Address.sin_port = htons(portNum);
    local_Address.sin_addr.s_addr = htonl(INADDR_ANY);

```

```

#pragma convert (819)
greeting = "This is a message from the C socket server.";
#pragma convert (0)

/* allocate socket */
if((server = socket(AF_INET, SOCK_STREAM, 0))<0)
{
    printf("failure on socket allocation\n");
    perror(NULL);
    exit(-1);
}

/* do bind */
if((bndrc=bind(server,(struct sockaddr*)&local_Address, address_len))<0)
{
    printf("Bind failed\n");
    perror(NULL);
    exit(-1);
}

/* invoke listen */
listen(server, 1);

/* wait for client request */
if((client = accept(server,(struct sockaddr*)NULL, 0))<0)
{
    printf("accept failed\n");
    perror(NULL);
    exit(-1);
}

/* send greeting to client */
if((sendrc = send(client, greeting, strlen(greeting),0))<0)
{
    printf("Send failed\n");
    perror(NULL);
    exit(-1);
}

close(client);
close(server);
}

```

For more information, see [Use sockets for interprocess communication](#).

Use input and output streams for interprocess communication

Input and output streams communicate between programs that are running in separate processes. The `java.lang.Runtime.exec()` method runs a program. The parent program can get handles to the child process input and output streams and can write to or read from those streams. If the child program is written in a language other than Java^(TM), you must ensure that any American Standard Code for Information Interchange (ASCII) or extended binary-coded decimal interchange code (EBCDIC) conversion takes place. See [Java character encodings](#) for more details.

For an example that uses input and output streams, see [Example: Use input and output streams for interprocess communication](#).

Example: Use input and output streams for interprocess communication: This example shows how to call a C program from Java^(TM) and use input and output streams for interprocess communication. In this example, the C program writes a string to its standard output stream, and the Java program reads this string and displays it. This example assumes that a library, which is named `JAVSAMPLIB`, has been created and that the `CSAMP1` program has been created in it.

Note: The JAVSAMPLIB is not created as part of the IBM Developer Kit licensed program (LP) number 5722-JV1 installation process. You must explicitly create it.

Example 1: CallPgm class

Note: Read the Code example disclaimer for important legal information.

```
import java.io.*;

public class CallPgm
{
    public static void main(String args[])
    {
        Process theProcess = null;
        BufferedReader inStream = null;

        System.out.println("CallPgm.main() invoked");

        // call the CSAMP1 program
        try
        {
            theProcess = Runtime.getRuntime().exec(
                "/QSYS.LIB/JAVSAMPLIB.LIB/CSAMP1.PGM");
        }
        catch(IOException e)
        {
            System.err.println("Error on exec() method");
            e.printStackTrace();
        }

        // read from the called program's standard output stream
        try
        {
            inStream = new BufferedReader(new InputStreamReader
                (theProcess.getInputStream()));
            System.out.println(inStream.readLine());
        }
        catch(IOException e)
        {
            System.err.println("Error on inStream.readLine()");
            e.printStackTrace();
        }
    }
} // end method

} // end class
```

Example 2: CSAMP1 C Program

Note: Read the Code example disclaimer for important legal information.

```
#include <stdio.h>
#include <stdlib.h>

void main(int argc, char* args[])
{
    /* Convert the string to ASCII at compile time */
    #pragma convert(819)
    printf("Program JAVSAMPLIB/CSAMP1 was invoked\n");
    #pragma convert(0)
    /* Stdout may be buffered, so flush the buffer */

    fflush(stdout);
}
```

For more information, see Use input and output streams for interprocess communication.

Example: Call Java from C

This is an example of a C program that uses the `system()` function to call the Java Hello program.

Example: Call Java from C

Note: Read the Code example disclaimer for important legal information.

```
#include <stdlib.h>

int main(void)
{
    int result;

    /* The system function passes the given string to the CL command processor
    for processing. */

    result = system("JAVA CLASS('com.ibm.as400.system.Hello')");
}
```

Example: Call Java from RPG

This is an example of an RPG program that uses the QCMDEXC API to call the Java^(TM) Hello program.

Example 1: Call Java from RPG

Note: Read the Code example disclaimer for important legal information.

```
D*           DEFINE THE PARAMETERS FOR THE QCMDEXC API
D*
DCMDSTRING   S           25   INZ('JAVA CLASS(''com.ibm.as400.system.Hello'')')
DCMDLENGTH   S           15P 5 INZ(25)
D*           NOW THE CALL TO QCMDEXC WITH THE 'JAVA' CL COMMAND
C           CALL 'QCMDEXC'
C           PARM                                CMDSTRING
C           PARM                                CMDLENGTH
C*         This next line displays 'DID IT' after you exit the
C*         Java Shell via F3 or F12.
C         'DID IT'   DSPLY
C*         Set On LR to exit the RPG program
C           SETON                                LR
C
```

Java platform

➤ The **Java^(TM) platform** is the environment for developing and managing Java applets and applications. It consists of three primary components: the Java language, the Java packages, and the Java virtual machine. The Java language and packages are similar to C++ and its class libraries. The Java packages contain classes, which are available in any compliant Java implementation. The application programming interface (API) should be the same on any system that supports Java.

Java differs from a traditional language, like C++, in the way it compiles and runs. In a traditional programming environment, you write and compile source code of a program into object code for a specific hardware and operating system. The object code binds to other object code modules to create a running program. The code is specific for a particular set of computer hardware and does not run on other systems without being changed. This figure illustrates the traditional language deployment environment.

To effectively use the Java platform, see the following:

Java applets and applications

You can write your Java applet and include it in an HTML page, much in the same way an image is included. When you use a Java-enabled browser to view an HTML page that contains an applet, the

applet's code is transferred to your system and is run by the browser's Java virtual machine. You can also write a Java application that does not require the use of a Web browser.

Java virtual machine

You can embed the Java virtual machine within a Web browser or in an operating system such as IBM^(R) Operating System/400^(R) (OS/400^(R)). The Java virtual machine consists of the Java interpreter and the Java runtime environment. The interpreter performs the task of interpreting the class file and running the Java instructions on a particular hardware platform. The Java virtual machine is what allows Java code to be written and compiled once, and run on any platform.

Java JAR and class files


The Java environment differs from other programming environments in that the Java compiler does not generate machine code for a hardware-specific instruction set. Instead, the Java compiler converts Java source code into Java virtual machine instructions, which Java class files store. You can use JAR files to store class files. The class file does not target a specific hardware platform, but instead targets the Java virtual machine architecture.

Java threads

Java is a multithreaded programming language; thus, more than one thread may be running within the Java virtual machine at one time. Java threads provide a way for a Java program to perform multiple tasks at the same time.


Java Development Kit

The Java Development Kit (JDK) is software that is distributed by Sun Microsystems, Inc. for Java developers. It includes the Java interpreter, Java classes, and Java development tools. Find the following information about JDKs:

- Java packages
- Java tools 

Java applets and applications

An **applet** is a Java^(TM) program designed to be included in an HTML Web document. The HTML document contains tags, which specify the name of the Java applet and its Uniform Resource Locator (URL). The URL is the location at which the applet bytecodes reside on the Internet. When an HTML document containing a Java applet tag is displayed, a Java-enabled Web browser downloads the Java bytecodes from the Internet and uses the Java virtual machine to process the code from within the Web document. These Java applets are what enable Web pages to contain animated graphics or interactive content.

For more information, see [Writing Applets](#) , Sun Microsystems' tutorial for Java applets. It includes an overview of applets, directions for writing applets, and some common applet problems.

Applications are stand-alone programs that do not require the use of a browser. Java applications run by starting the Java interpreter from the command line and by specifying the file that contains the compiled application. Applications usually reside on the system on which they are deployed. Applications access resources on the system, and are restricted by the Java security model.

Java virtual machine

The Java^(TM) virtual machine is a runtime environment that you can add into a web browser or any operating system, such as IBM Operating System/400 (OS/400). The Java virtual machine runs instructions that a Java compiler generates. It consists of a bytecode interpreter and runtime that allow Java class files to run on any platform, regardless of the platform on which they were originally developed.

The class loader and security manager, which is part of the Java runtime, insulate code that comes from another platform. They also can restrict which system resources each class that is loaded accesses.

Note: Java applications are not restricted; only applets are restricted. Applications can freely access system resources and use native methods. Most IBM Developer Kit for Java programs are applications.

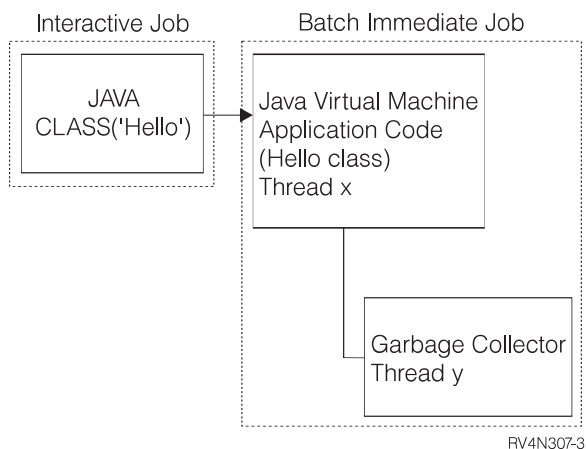
You can use the Create Java Program (CRTJVAPGM) command to ensure that the code meets the safety requirements that the Java runtime imposes to verify the bytecodes. This includes enforcing type restrictions, checking data conversions, ensuring that parameter stack overflows or underflows do not occur, and checking for access violations. However, you do not need to explicitly verify the bytecodes. If you do not use the CRTJVAPGM command in advance, then the checks occur during the first use of a class. Once the bytecodes are verified, the interpreter decodes the bytecodes and runs the machine instructions that are needed to carry out the desired operations.

Note: The Java interpreter is only used if you specify OPTIMIZE(*INTERPRET) or INTERPRET(*YES).

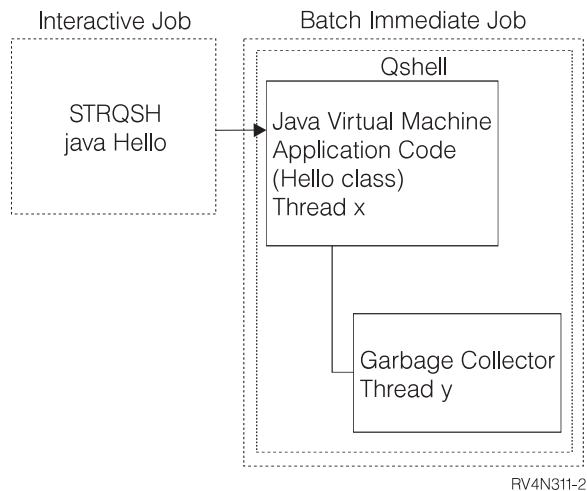
In addition to loading and running the bytecodes, the Java virtual machine includes a garbage collector that manages memory. Garbage collection runs at the same time as the loading and interpretation of the bytecodes.

Java runtime environment

The Java runtime environment starts whenever you enter the Run Java (RUNJVA) command or JAVA command on the iSeries command line. Because the Java environment is multithreaded, it is necessary to run the Java virtual machine in a job that supports threads, such as a batch immediate (BCI) job. Once the Java virtual machine starts, additional threads may start in which the garbage collector runs. The typical Java environment looks like this:



It is also possible to start the Java runtime environment by using the java command in Qshell from the Qshell Interpreter. In this environment, the Qshell Interpreter is running in a BCI job that is associated with an interactive job. The Java runtime environment starts in the job that is running the Qshell Interpreter.



RV4N311-2

When the Java runtime environment starts from an interactive job, the Java Shell Display is shown. This display provides an input line for entering data into the System.in stream, as well as displaying data that is written to the System.out stream and System.err stream.

Java interpreter

The Java interpreter is the part of the Java virtual machine that interprets Java class files for a particular hardware platform. The Java interpreter decodes each bytecode and runs a series of machine instructions for that bytecode.

Java JAR and class files

A **Java[™] ARchive (JAR) file** is a file format that combines many files into one. You can use JAR as a general archiving tool and also to distribute Java programs of all types, including applets. Java applets download into a browser in a single Hypertext Transfer Protocol (HTTP) transaction rather than by opening a new connection for each piece. This method of downloading improves the speed at which an applet loads on a Web page and begins functioning.

JAR is the only archive format that is cross-platform. JAR is also the only format that handles audio files and image files, as well as class files. JAR is an open standard, fully extendable format that is written in Java.

The JAR format also supports compression, which reduces the size of the file and decreases download time. Additionally, an applet author may digitally sign individual entries in a JAR file to authenticate their origin.

➤ To update classes in JAR files, see the Java jar tool. ⏪

Java class files are stream files that are produced when a source file is compiled by the Java compiler. The class file contains tables that describe each field and method of the class. The file also contains the bytecodes for each method, static data, and descriptions that are used to represent Java objects.

Java threads

A **thread** is a single independent stream that runs within a program. Java[™] is a multithreaded programming language, so more than one thread may be running within the Java virtual machine at one time. Java threads provide a way for a Java program to perform multiple tasks at the same time. A thread is essentially a flow of control in a program.

Threads are a modern programming construct that are used to support concurrent programs and to improve the performance and scalability of applications. Most programming languages support threads through the use of add-in programming libraries. Java supports threads as built-in application program interfaces (APIs).

Note: The use of threads provides the support to increase the interactivity, meaning a shorter wait at the keyboard because more tasks are running in parallel. But, the program is not necessarily more interactive just because it has threads.

Threads are the mechanism for waiting on long running interactions, while still allowing the program to handle other work. Threads have the ability to support multiple flows through the same code stream. They are sometimes called **lightweight processes**. The Java language includes direct support for threads. But, by design, it does not support asynchronous non-blocking input and output with interrupts or multiple wait.

Threads allow the development of parallel programs that scale well in an environment where a machine has multiple processors. If properly constructed, they also provide a model for handling multiple transactions and users.

You can use threads in a Java program for a number of situations. Some programs must be able to engage in multiple activities and still be able to respond to additional input from the user. For example, a Web browser should be able to respond to user input while playing a sound.

Threads can also use asynchronous methods. When you call a second method, you do not have to wait for the first method to complete before the second method continues with its own activity.

There are also many reasons not to use threads. If a program uses inherently sequential logic, one thread can accomplish the entire sequence. Using multiple threads in such a case results in a complex program with no benefits. There is considerable work in creating and starting a thread. If an operation involves only a few statements, it is faster to handle it in a single thread. This can be true even when the operation is conceptually asynchronous. When multiple threads share objects, the objects must synchronize to coordinate thread access and maintain consistency. Synchronization adds complexity to a program, is difficult to tune for optimal performance, and can be a source of programming errors.

For more threads information, see Developing multithreaded applications.

Sun Microsystems, Inc. Java Development Kit


The Java^(TM) Development Kit (JDK) is software that is distributed by Sun Microsystems, Inc. for Java developers. It includes the Java interpreter, Java classes, and Java development tools: compiler, debugger, disassembler, appletviewer, stub file generator, and documentation generator.

The JDK enables you to write applications that are developed once and run anywhere on any Java virtual machine. Java applications that are developed with the JDK on one system can be used on another system without changing or recompiling the code. The Java class files are portable to any standard Java virtual machine.

To find more information about the current JDK, check the version of the IBM Developer Kit for Java on your iSeries server.

You can check the version of the default IBM Developer Kit for Java Java virtual machine on your iSeries server by entering either of the following commands:

- `java -version` on the Qshell command prompt.
- `RUNJAVA CLASS(*VERSION)` on the CL command line.

Then, look for the same version of Sun Microsystems, Inc. JDK at The Source for Java Technology java.sun.com  for specific documentation. The IBM Developer Kit for Java is a compatible implementation of the Sun Microsystems, Inc. Java Technology, so you should be familiar with their JDK documentation.

See the following topics for more information:

- Support for multiple Java Development Kits (JDKs) provides information about using different Java virtual machines.
- Native methods and the Java Native Interface defines what a native method is and what they do. This topic also briefly explains the Java Native Interface.

Java packages

A Java package is a way of grouping related classes and interfaces in Java. Java packages are similar to class libraries that are available in other languages.

The Java packages, which provide the Java APIs, are available as part of Sun Microsystems, Inc. Java Development Kit (JDK).

Package	Contents
java.applet	Applet classes
java.awt	Graphics, window, and graphical user interface (GUI) classes
java.awt.datatransfer	Data transfer classes
java.awt.event	Event processing classes and interfaces
java.awt.image	Image processing classes
java.awt.peer	GUI interfaces for platform independence
java.beans	JavaBeans component model API
java.io	Input and output classes
java.lang	Core language classes
java.lang.reflect	Reflection API classes
java.math	Arbitrary precision arithmetic
java.net	Networking classes
java.rmi	Remote Method Invocation classes
java.rmi.dgc	RMI-related classes
java.rmi.registry	RMI-related classes
java.rmi.server	RMI-related classes
java.security	Security classes
java.security.acl	Security-related classes
java.security.interfaces	Security-related classes
java.sql	JDBC SQL API for database classes
java.text	Internationalization classes
java.util	Data types
java.util.zip	Compression and decompression classes

For more information about the Java APIs by Sun Microsystems, Inc., see the Sun Microsystems, Inc. API User's Guide.

Java tools

For a complete list of tools that Sun Microsystems, Inc. Java Development Kit supplies, see Tools Reference by Sun Microsystems, Inc. For more information about each individual tool that the IBM Developer Kit for Java supports, see Java tools that are supported by the IBM Developer Kit for Java.

Advanced topics

» The following are advanced topics for the IBM Developer Kit for Java^(TM):

Classes, packages, and directories

Each Java class is part of a package. The package name relates to the directory structure in which the class resides.

Files in the integrated file system

The integrated file system stores Java-related class, source, ZIP, and JAR files in a hierarchical file structure.

File authorities

To run or debug a Java program, your class, ZIP, or JAR files require read authority. Find more information about file authorities that several CL commands require.

Batch job

Java programs can run in a batch job by using the Submit Job (SBMJOB) command. Find out more about the SBJJOB command and how you can verify that your batch job is capable of running more than one job.



Java classes, packages, and directories

Each Java^(TM) class is part of a package. The first statement in a Java source file indicates which class is in what package. If the source file does not contain a package statement, the class is part of an unnamed default package.

The package name relates to the directory structure in which the class resides. The integrated file system supports Java classes in a hierarchical file structure that is similar to what you find on most PC and UNIX systems. You must store a Java class in a directory with a relative directory path that matches the package name for that class. For example, consider the following Java class:

```
package classes.geometry;
import java.awt.Dimension;

public class Shape {

    Dimension metrics;

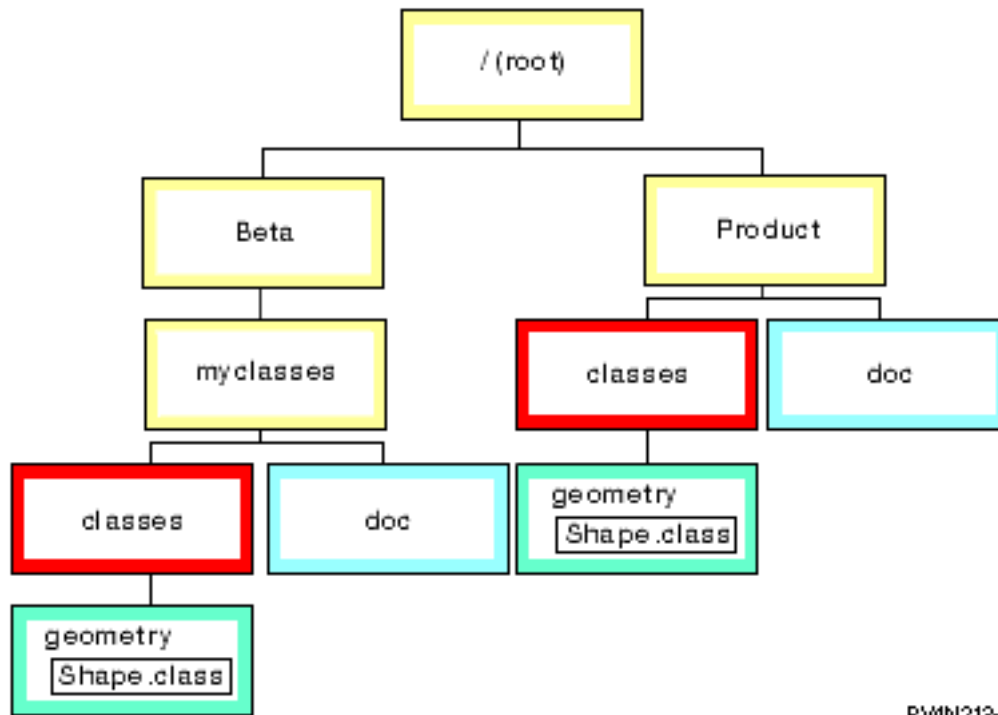
    // The implementation for the Shape class would be coded here ...

}
```

The package statement in the previous code indicates that the Shape class is part of the classes.geometry package. For the Java runtime to find the Shape class, store the Shape class in the relative directory structure classes/geometry.

Note: The package name corresponds to the relative directory name in which the class is stored. The Java virtual machine class loader finds the class by appending the relative path name to each directory that is specified in the classpath. The Java virtual machine class loader can also find the class by searching the ZIP files or JAR files that are specified in the classpath.

For example, if the Shape class was stored in the /Product/classes/geometry directory in the "root" (/) file system, /Product would need to be specified in the classpath.



RV4N312-1

Note: Multiple versions of the Shape class can exist in the directory structure. To use the "beta" version of the Shape class, place /Beta/myclasses in the CLASSPATH before any other directories or ZIP files that contain the Shape class.

The Java compiler uses the Java classpath, package name, and directory structure to find packages and classes when compiling Java source code.

Files in the integrated file system

The integrated file system stores Java-related class files, source files, ZIP files, and JAR files in a hierarchical file structure. You can also store source files in the integrated file system. You may store the files in these integrated file systems:

- "Root" (/) file system
- Open systems file system (QOpenSys)
- User-defined file system
- Library file system (QSYS.LIB)
- OS/2 Warp Server for iSeries file system (QLANSrv)
- Optical file system (QOPT)

Note: Other integrated file systems are not supported because they are not threadsafe.

Java file authorities in the integrated file system

To run or debug a Java[™] program, the class file, JAR file, or ZIP file needs to have read authority (*R). Any directories need read and execute authorities (*RX).

To use the Create Java Program (CRTJVAPGM) command to optimize a program, the class file, JAR file, or ZIP file must have read authority (*R), and the directory must have execute authority (*X). If you use a pattern in the class file name, the directory must have read and execute authority (*RX).

To delete a Java program by using the Delete Java Program (DLTJVAPGM) command, you must have read and write authority (*RW) to the class file, and the directory must have execute authority (*X). If you use a pattern in the class file name, the directory must have read and execute authority (*RX).

To display a Java program by using the Display Java Program (DSPJVAPGM) command, you must have read authority (*R) to the class file, and the directory must have execute authority (*X).

Note: Files and directories that do not have execute authority (*X) always appear to have execute authority (*X) to a user with QSECOFR authority. Different users can get different results in certain situations, even though both users appear to have the same access to the same files. This is important to know when running shell scripts using the Qshell Interpreter or `java.Runtime.exec()`.

For example, one user writes a Java program that uses `java.Runtime.exec()` to call a shell script, then tests it using a user ID with QSECOFR authority. If the file mode of the shell script has read and write authority (*RW), the integrated file system allows the user ID with QSECOFR authority to run it. However, a non-QSECOFR authority user could try to run the same Java program, and the integrated file system would tell the `java.Runtime.exec()` code that the shell script cannot be run, because *X is missing. In this case, `java.Runtime.exec()` throws an input and output exception.

➤ You can also assign authorities to new files created by Java programs in an integrated file system. By using the `os400.file.create.auth` system property for files and `os400.dir.create.auth` for directories, any combination of read, write, and execute authorities may be used. ⬅

For more information, see the Program and CL Command APIs or the Integrated file system.

Run Java in a batch job

Java^(TM) programs run in a batch job by using the Submit Job (SBMJOB) command. In this mode, the Java Qshell Command Entry display is not available to handle the `System.in`, `System.out`, nor `System.err` streams.

You may redirect these streams to other files. Default handling sends the `System.out` and `System.err` streams to a spooled file. The batch job, which results in an input and output exception for read requests from `System.in`, owns the spooled file. You can redirect `System.in`, `System.out`, and `System.err` within your Java program. You can also use the `os400.stdin`, `os400.stdout`, and `os400.stderr` system properties to redirect `System.in`, `System.out`, and `System.err`.

Note: SBJJOB sets the Current Working Directory (CWD) to the HOME directory that is specified in the user profile.

Example: Running Java in a Batch Job

```
SBMJOB CMD(JAVA He11o OPTION(*VERBOSE)) CPYENVVAR(*YES)
```

Running the `JAVA` command in the previous example spawns a second job. Therefore, the subsystem that the batch job runs in must be capable of running more than one job.

You can verify that your batch job is capable of running more than one job by following these steps:

1. On the CL command line, enter `DSPSBSD(MYSBSD)`, where `MYSBSD` is the subsystem description of your batch job.
2. Choose option 6, Job queue entries.

3. Look at the Max Active field for your job queue.

If the Max Active field is less than or equal to 1, and it is not *NOMAX, enter the following on the CL command line:

```
CHGJOBQE SBSD(MYSBSD) JOBQ(MYJOBQ) MAXACT(*NOMAX)
```

Where:

- MYSBSD is your subsystem description, and
- MYJOBQ is your job queue.

Run your Java application on a host that does not have a graphical user interface

If you want to run your Java^(TM) application on a host that does not have a graphical user interface (GUI), such as an iSeries server, you can use the Remote Abstract Window Toolkit (AWT), the Class Broker for Java (CBJ) or the Native Abstract Windowing Toolkit (NAWT).

You use Remote AWT with install and administration interfaces of server applications. These interfaces typically have a minimum of complex graphics and highly-interactive content. Remote AWT distributes AWT processing between the iSeries server and a workstation. So, responsiveness of graphic-intensive and highly-interactive operations are not as fast as AWT implementations on platforms with locally-attached graphic terminals. To use Remote AWT, see set up Remote AWT.

You can use the CBJ for high performance GUI services. Since Remote AWT is not recommended for complex graphics or highly-interactive operations, you can use the CBJ instead, which was designed for these environments. To use CBJ, see set up CBJ.

➤ You can use NAWT for Java graphical computation entirely on an iSeries server using the X Window System. The **X Window System** is a graphics system that provides a client/server base for displaying graphics. The X Window graphics server is highly portable, allowing support for a variety of languages and operating systems. NAWT provides Java applications and servlets with the capability to use the Java Development Kit's (JDK) AWT graphics functionality. ⚡

IBM Developer Kit for Java Remote Abstract Window Toolkit

The Remote Abstract Window Toolkit is an implementation of the Abstract Window Toolkit (AWT). It allows Java^(TM) applications to run, without any changes, on a host that does not have a graphical user interface (GUI). An iSeries server does not support locally-attached graphic terminals; therefore, Remote AWT is necessary to allow graphical Java applications to run on the iSeries server.

You use Remote AWT with installation and administration interfaces of server applications. These interfaces typically have a minimum of complex graphics and highly-interactive content. Remote AWT distributes AWT processing between an iSeries server and a workstation. Responsiveness of graphic-intensive and highly-interactive operations may not be as fast as AWT implementations on platforms with locally-attached graphic terminals.

Since IBM Developer Kit for Java Remote AWT is not recommended for complex graphics or highly-interactive operations, you can use the Class Broker for Java instead, which was designed for these environments.

For information on how to set up Remote AWT, see Set up the Remote Abstract Window Toolkit for Java on a remote display.

Set up the Remote Abstract Window Toolkit for Java on a remote display

With the Remote Abstract Window Toolkit (AWT), you can run Java^(TM) AWT graphics programs without making any program source changes and display the graphics remotely. To use Remote AWT you must

have Transmission Control Protocol/Internet Protocol (TCP/IP) set up, and Sun Microsystems, Inc., Java Development Kit (JDK) 1.1.8 or Java 2 SDK (J2SDK), Standard Edition, installed on your iSeries server and remote display.

You can use any graphics-capable hardware, including IBM Network Station, as a remote display for Remote AWT if it meets these requirements:


- Graphics-capable hardware that runs Windows^(R) 95, Windows NT 4.0, IBM Operating System/2^(R) (OS/2^(R)), Sun Solaris, or AIX^(R)
- Configured hardware to access your iSeries server with TCP/IP
- Java Development Kit 1.1.8 or J2SDK

To set up Remote AWT, complete these tasks:

1. Make the Remote AWT class files accessible to the remote display by copying the files to the remote display or mapping the path to a network drive on the remote display.
2. Add RAWTGui.zip or RAWTGui.jar to the CLASSPATH of the remote display. For JDK 1.1.8, add the RAWTGui.zip file to the CLASSPATH of the remote display by setting the CLASSPATH environment variable or by using the `-classpath` parameter of the `java` command. For J2SDK, the RAWTGui.jar file is automatically added to the CLASSPATH when using `-jar` parameter of the `java` command.
3. Start Remote AWT on the remote display.

For details and hints for using Remote AWT, see these topics:

- Run a Java program using Remote Abstract Window Toolkit provides instructions on how to run a Java program with an iSeries server using multiple JDKs and Netscape.
- Print with the Remote Abstract Window Toolkit explains how to print, which is the same as standard Java AWT printing. It also shows you how to print to your iSeries server.
- Remote Abstract Window Toolkit properties shows you how to run a Remote AWT application using the `os400.class.path.rawt` property.
- Remote Abstract Window Toolkit SecurityManager restrictions provides information about the restrictions that apply when you are running Java applications using Remote AWT under the control of a SecurityManager.

For more information about setting up TCP/IP, see "How do I set up TCP/IP" in the TCP/IP Configuration and Reference, SC41-5420 book. 

See Example: Set up the Remote Abstract Window Toolkit for Java on a Windows remote display, for an example of how to set up Remote AWT.

For more information about AWT, see the Abstract Window Toolkit by Sun Microsystems, Inc.

Make the Remote Abstract Window Toolkit for Java class files accessible to the remote display

To make the Remote Abstract Window Toolkit (AWT) class files accessible to the remote display, follow these steps for either Java^(TM) Development Kit (JDK) 1.1. x or Java 2 SDK (J2SDK), Standard Edition, version 1.2. It is important to note, however, that in order for the Remote AWT to function correctly, the version of the RAWTGui.jar file you use on the graphics-capable remote display must match the JDK or J2SDK version you use on the host.

If you are using JDK 1.1.8, you can either:

- Copy the Remote AWT class files to the remote display.
The Remote AWT files are installed with the IBM Developer Kit for Java in two ZIP files:
 - /QIBM/ProdData/Java400/jdk118/RAWTApplHost.zip
 - /QIBM/ProdData/Java400/jdk118/RAWTGui.zip

The RAWTAppHost.zip file contains the Remote AWT classes for the iSeries server. The RAWTGui.zip file contains the Remote AWT classes for your remote display.

Copy RAWTGui.zip from /QIBM/ProdData/Java400/jdk118 to your remote display.

- Map the path, /QIBM/ProdData/Java400/jdk118/RAWTGui.zip, to a network drive on your remote display.

If you are using J2SDK, version 1.2 or higher, you can either:

- Copy the Remote AWT class files to the remote display.

The Remote AWT files are installed with the IBM Developer Kit for Java in two JAR files:

- /QIBM/ProdData/Java400/jdk12/RAWTAHost.jar
- /QIBM/ProdData/Java400/jdk12/RAWTGui.jar

If you are using a version of the J2SDK other than 1.2, substitute that version number into all path instances in this section.

The RAWTAHost.jar file contains the Remote AWT classes for the iSeries server. The RAWTGui.jar file contains the Remote AWT classes for your remote display.

Copy RAWTGui.jar from /QIBM/ProdData/Java400/jdk12 to a network drive on your remote display.

- Map the path, /QIBM/ProdData/Java400/jdk12/RAWTGui.jar, to a network drive on your remote display.

Add RAWTGui.zip or RAWTGui.jar to the CLASSPATH of the remote display

Setting the CLASSPATH allows the Remote Abstract Window Toolkit (AWT) classes to be found by the Java^(TM) virtual machine on the remote display. This step is only necessary for JDK 1.1.x. This step is not necessary with Java 2 SDK (J2SDK). To add the RAWTGui.zip file to the CLASSPATH of the remote display, do either of these steps:

- Set the CLASSPATH environment variable. See the Java Development Kit (JDK) information for your remote display for details.

Add the path where the RAWTGui.zip file is located to the CLASSPATH environment variable.

- Use the -classpath parameter of the java command.

When you use the java command to start Remote AWT, you can use the -classpath parameter to specify a CLASSPATH. The CLASSPATH includes the path where the RAWTGui.zip file is located.

For example, in Windows^(R), the CLASSPATH parameter may look like this:

```
-classpath c:\jdk1.1.7\lib\classes.zip;c:\rawt\RAWTGui.zip
```

J2SDK, version 1.2 and higher JAR support sets the CLASSPATH, so you do not need to explicitly set the CLASSPATH parameter. To set the classpath and start Remote AWT on a remote display, enter this command:

```
java -jar <PATH>RAWTGui.jar
```

where <PATH> is the fully qualified drive and directory where the RAWTGui.jar file is located. For example, java -jar c:\rawt2\RAWTGui.jar.

Start the Remote Abstract Window Toolkit for Java on a remote display

You need to start the server daemon on your remote display once and it stays active until you end it. Java^(TM) programs that exit on an iSeries server do not end the server daemon.

Note: The Welcome dialog stays active when you start the server daemon. When the Welcome dialog display closes, the server daemon ends. You can minimize the Welcome dialog display while the server daemon is active and use the display to end the server daemon.

To start the Remote Abstract Window Toolkit (AWT) server daemon for JDK 1.1.x, enter this on the command line:

```
java -classpath <PATH>RAWTGui.zip;C:\jdk1.1.8\lib\classes.zip
com.ibm.rawt.server.RAWTPCServer
```


where <PATH> is the fully qualified drive and directory where the RAWTGui.jar file is located. For example, java -jar c:\rawt2\RAWTGui.jar.

To start the Remote AWT server daemon for J2SDK, version 1.3, enter this on the command line:

```
java -jar <PATH>RAWTGui.jar
```

where <PATH> is the fully qualified drive and directory where the RAWTGui.jar file is located. For example, java -jar c:\rawt2\RAWTGui.jar.

The server daemon selects the first free port above 2000 when the Java application connects using Remote AWT. The Java application uses this port until the application ends. Additional Java applications are connected to subsequent free ports above 2000. The available range of ports goes up to 9999.

For more information about setting up TCP/IP, see "How do I set up TCP/IP" in the TCP/IP Configuration and Reference, SC41-5420 book. 

Run a Java program using Remote Abstract Window Toolkit

To run Java^(TM) using the Remote Abstract Window Toolkit (AWT), perform the following steps:

1. Start Remote AWT on the remote display.
2. Start the Java program on the iSeries server.
 - a. Enter the Run Java (RUNJAVA) command on the command line.
Note: You must define the Java classpath to the Java program.
 - b. Press F4 (Prompt).
 - c. Enter the Java program class name on the class parameter line.
 - d. Press F10 (Additional parameters).
 - e. Press the Page Down key.
 - f. Enter RmtAwtServer on the next property name parameter line.
 - g. Enter the Transmission Control Protocol/Internet Protocol (TCP/IP) address (for example, 1.1.11.11) of the remote display on the next property value parameter line.
 - h. Enter os400.class.path.rawt on the property name parameter line.
 - i. Enter 1 on the property value parameter line.
 - j. Enter + for more properties.
 - k. Enter java.version on the property name parameter line.
 - l. Enter 1.3 on the property value parameter line. This version must match the version of the RAWTGui.jar daemon code running on the remote display.

The command line should have the following pattern, all on one line:

```
java class (classname) prop(('RmtAwtServer' '1.1.11.11')
('os400.class.path.rawt' '1')('java.version' '1.3'))
```

- m. Press Enter.

You can also run a Java program using Remote AWT with Netscape.

Run a Java program using Remote Abstract Window Toolkit with Netscape: When running Java^(TM) applications with Netscape, you can run them one of two ways.

One option is to start the Remote Abstract Window Toolkit (AWT) server within the Netscape Java virtual machine by opening an HTML file that contains `com.ibm.rawt.server.StartRAWT.class`. For example, see the `RAWT.html` file below. Once started, you can then start a Java application on your iSeries server.

Or, you can start the Remote AWT server within the Netscape Java virtual machine by opening an HTML file that contains `com.ibm.rawt.server.StartRAWT400.class` and the IBM Toolbox for Java classes. For example, see the `RAWT400.html` file shown below. Once started, you can sign on to your iSeries server where the Java application resides and start the application.

Running within the Remote AWT server within the Netscape Java virtual machine

To run the Remote AWT server within the Netscape Java virtual machine, follow these steps:

1. Edit this example `.html` file for your specific installation information for `RAWTGui.zip`. This file, `RAWT.html`, starts Remote AWT within the Netscape Java virtual machine.

Example: Start Remote AWT within the Netscape Java virtual machine

Note: Read the Code example disclaimer for important legal information.

```
<HTML>
<BODY TEXT="#000000" LINK="#0000EE" VLINK="#551A8B" ALINK="#FF0000">
<CENTER>
<APPLET CODE="com.ibm.rawt.server.StartRAWT.class"
codebase="file://C|remote_aws\jdk1.1.7\lib\RAWTGui.zip"
WIDTH=600 HEIGHT=50>
</APPLET>
</CENTER>
</BODY>
</HTML>
```

2. Browse the `RAWT.html` page with Netscape 4.05 or higher. After granting all of the requested privileges, Netscape starts the Remote AWT server and runs it within its Java virtual machine.
3. Start a Java application on your iSeries server using Remote AWT.
4. After you exit the application, click on the Reload button while pushing the Shift key to start the Remote AWT server again.

Running the Remote AWT server within the Netscape Java virtual machine and signing on to your iSeries server

To run the Remote AWT server within the Netscape Java virtual machine and sign on to your iSeries server, follow these steps:

1. Edit this example `.html` file for your specific installation information for `jt400.zip` and `RAWTGui.zip`. This file, `RAWT400.html`, starts Remote AWT and uses the IBM Toolbox for Java to sign on to the iSeries server.

Example: Start Remote AWT and use the Toolbox for Java to sign onto an iSeries server

Note: Read the Code example disclaimer for important legal information.

```
<HTML>
<BODY TEXT="#000000" LINK="#0000EE" VLINK="#551A8B" ALINK="#FF0000">
<CENTER>
<APPLET ARCHIVE="file://C|jt400\lib\jt400.zip"
code="com.ibm.rawt.server.StartRAWT400.class"
codebase="file://C|remote_aws\jdk1.1.1\lib\RAWTGui.zip"
WIDTH=600 HEIGHT=50>
</APPLET>
</CENTER>
</BODY>
</HTML>
```

2. Browse this `RAWT400.html` page with Netscape 4.05. After granting all of the requested privileges, Netscape starts the Remote AWT applet that displays a panel where you can do any of these options:

- Sign on to your iSeries server with Remote AWT using IBM Toolbox for Java to access the iSeries server.
- Enter the Java application name and arguments with Remote AWT properties.
- Push the Start Application button to start the specified Java application with Remote AWT.

Print with the Remote Abstract Window Toolkit

Printing with the Remote Abstract Window Toolkit (AWT) is the same as standard Java[™] AWT printing. The Remote AWT remote display processes print output and directs the output to any printer that is known to the remote display operating system. This can be a printer that is directly attached to the remote display, or a network printer that is known to the remote display operating system.

You can select to print on the remote display or to an iSeries server. A new print dialog displays when an application issues a print request. The print request allows you to select either a Remote Display printer or OS/400 printer. If you select OS/400 printer, a sign on dialog display appears. Once you have signed on, the Print Dialog display appears. You can specify the OS/400 print queue, print file, file and banner page title. You can also select paper size, orientation, and number of copies.

To use remote printing, you must install the IBM Toolbox for Java (5763-JC1) and add this to your classpath on the iSeries server:

```
QIBM/ProdData/HTTP/Public/jt400/lib/jt400.zip
```

You can update the classpath by adding the classpath environment variable or by using the classpath parameter.

Note: If this message appears while printing to the iSeries server, either the IBM Toolbox for Java is not installed or the IBM Toolbox for Java classes are not in the classpath.

```
Failed to load class file: com/ibm/as400/access/PrintObjectList.class
Exception occurred during event dispatching:
java.lang.NoClassDefFoundError: com/ibm/as400/access/PrintObjectList
```

Remote Abstract Window Toolkit properties

When running a Java[™] Remote AWT application on the iSeries server, the `os400.class.path.rawt` property must be used with a value of 1. There are a number of default properties that are required when you use Remote AWT. These default properties along with the proper version of Remote AWT and the CLASSPATH are set when you use the `os400.class.path.rawt` property. The version of Remote AWT is set according to the JDK version, which is either the default version, if not specified, or the version that you specify with the `java.version` property.

Here are the default properties that are required for Remote AWT.

For JDK 1.1.x:

- `awt.toolkit=com.ibm.rawt.CToolkit`

For J2SDK:

- `awt.toolkit=com.ibm.rawt2.ahost.java.awt.AHToolkit`
- `java.awt.graphicsenv=com.ibm.rawt2.ahost.java.awt.AHGraphicsEnvironment`
- `java.awt.printerjob=com.ibm.rawt2.ahost.java.awt.print.AHPrinterjob`

Remote AWT properties for remote display

If the server daemon or Java application abnormally ends with this message, check the Java version on your remote display.

```
The JDK version in the Application-host/User-station is incompatible
with the Remote-AWT version...
```

To check the version level, enter `java -version` on the command line. You can use this new property on a remote display if you have a problem with the JDK version. This property is not applicable for the iSeries server. If the version is not at level 1.1.x, then you must install the proper level. If it is at level 1.1.x, you can run the Remote AWT server or the Java application with this property that indicates the Java version:

```
-DJdkVersion=1.1.x
```

Remote Abstract Window Toolkit SecurityManager restrictions

Java^(TM) SecurityManager is not commonly used. If a SecurityManager is installed, however, it must allow the listed calls to succeed for RAWT to operate.

- SecurityManager.checkAccess(..)
- SecurityManager.checkMemberAccess(..)
- SecurityManager.checkExit(..)
- SecurityManager.checkRead(String file)
- SecurityManager.checkConnect(...)
- SecurityManager.checkListen(...)
- SecurityManager.checkAccept(...)
- SecurityManager.checkPropertiesAccess(..)

Example: Set up the Remote Abstract Window Toolkit for Java^(TM) on a Windows remote display

This example shows one way that you can set up Remote AWT on a Windows^(R) remote display. There are many other ways you can do this, depending on your preference. You can use a similar process on other remote display operating systems. The setup and starting processes are automated with a Windows .bat file or other programming facility that the remote display operating system provides.

To set up the Remote Abstract Window Toolkit (AWT) for Java^(TM) on a Windows remote display, do these tasks:

- Make the Remote AWT class files accessible to the remote display.
Copy the Remote AWT class files to the remote display. Perform one of the following tasks:
 - Copy `/QIBM/ProdData/Java400/jdk118/RAWTGui.zip` to `c:\rawt\RAWTGui.zip`.
 - Copy `/QIBM/ProdData/Java400/jdk1x/RAWTGui.jar`, where `x` is your Java 2 Software Development Kit (J2SDK), Standard Edition version (either 2, 3, or 4), to `c:\rawt2\RAWTGui.jar`.
- Start the Remote AWT on the remote display by entering this on the command line:

```
java -classpath c:\jdk1.1.8\lib\classes.zip;c:\rawt\RAWTGui.zip  
java com.ibm.rawt.server.RAWTPCServer
```

or

```
java -jar c:\rawt2\RAWTGui.jar
```

For more information, see [Set up the Remote Abstract Window Toolkit for Java on a remote display](#).

Class Broker for Java

Class Broker for Java^(TM) (CBJ) is a general-purpose framework for writing client/server applications in Java. Typically, a client/server application consists of both a client object and a server object. The server and client are responsible for all communication between the two objects. This communication is often accomplished by Remote Method Invocation (RMI) or by a socket connection. RMI is not easy or flexible to use and there is a considerable learning curve when trying to use sockets effectively.

CBJ is easy to use, flexible, and hides the complexity of socket connections. Except for a few calls to the CBJ classes to initialize the application, a brokered client/server application appears to be a local application. CBJ handles all communication and resource loading between the client and server. The fact that some of the objects of the program are running on the client and some are running on the server is

almost transparent. CBJ creates client and server proxies using the CBJ runtime. Once the broker creates the proxy objects, the client communicates with the remote server by calling methods on the proxy of the server. Likewise, the server object communicates with the client by calling methods on the proxy of the client. Thus, both the client and server sides of the application appear as if they are calling methods of local objects.

For information about using Class Broker for Java, see [Set up Class Broker for Java](#).

Set up the Class Broker for Java on a remote display

With the Class Broker for Java^(TM) (CBJ), you can run Class Broker for Java-enabled Java graphical programs on the iSeries server and display the graphics remotely.

You can either install CBJ yourself or have the system administrator install it for you. If the system administrator installs the product, all programmers can share the same Java code.

You can either install CBJ on Windows^(R) 95/98/NT, UNIX^(R), or iSeries server. In most circumstances, you must install CBJ on your client machine and server machine.

Note: If your client machine is running a client applet that is accessed through a web server on the server machine, then you do not need to install CBJ on the client machine.

To install CBJ on the iSeries server, see [Install Class Broker for Java on an iSeries server](#).

To install CBJ on your workstation, see [Install Class Broker for Java on Windows or UNIX](#).

Install Class Broker for Java on an iSeries server

To install Class Broker for Java^(TM) (CBJ) on an iSeries server, follow these steps:

1. Make sure that the IBM Developer Kit for Java is properly installed. See [Install the IBM Developer Kit for Java](#) to test installation.
2. Select a directory in the integrated file system to store the CBJ package named `cbj_1.1.jar`. For example, `/usr/local`.
3. Start the Qshell Interpreter by entering the QSH command and go to the directory that you selected in the integrated file system. For example, `/usr/local`. When you install the IBM Developer Kit for Java, the `cbj_1.1.jar` package is installed in `QIBM/ProdData/Java400/ext`.
4. Enter this command in Qshell:

```
jar xvf "PATH"cbj_1.1.jar
```

"PATH" is the directory path where the `cbj_1.1.jar` package is located. For example, `QIBM/ProdData/Java400/ext`.

The CBJ files are extracted into the subdirectory, named `/usr/local/JCBroker`. See [Package content of cbj_1.1.jar](#) for more information.

5. Create an Java program on your iSeries server for `jcb.jar`, by entering this command:

```
CRTJVAPGM CLSF('/usr/local/JCBroker/lib/jcb.jar')
```

6. If you want to use the CBJ classes (not in debug mode), you have to add `JCBroker\lib` and `JCBroker\lib\jcb.jar` to the classpath option in the Java command line. It is not recommended that you add `JCBroker\lib` and `JCBroker\lib\jcb.jar` to the CLASSPATH environment variable, because it might conflict with class loading when you run applets that set `applet_jcb.jar` in the ARCHIVE tag.
7. If you want to run CBJ in debug mode, replace `jcb.jar` with `jcbd.jar` in the classpath or override the classpath value on the Java command line. You can also use this command:

```
CRTJVAPGM CLSF('/usr/local/JCBroker/lib/jcbd.jar')
```

8. For more information on CBJ APIs, running demos, editing properties, designing and writing CBJ applications, and other topics, refer to the `JCBroker/index.html` file in the `cpj_1.1.jar` package.

Install Class Broker for Java on Windows or UNIX

To install Class Broker for Java^(TM) (CBJ) on Windows^(R), follow these steps:

1. Make sure that either JDK/JRE1.1 or JDK/JRE1.2 is properly installed.
2. Select a directory in to store the CBJ package, named `cbj_1.1.jar`. For example, `C:\`.
3. Change to your directory (`C:\`) and enter this command:

```
C:\ > jar xvf cbj_1.1.jar
```

The CBJ files are extracted and are copied into the directory. See Package content of `cbj_1.1.jar` for more information.

4. If you want to use the CBJ classes (not in debug mode), add `C:\JCBroker\lib` and `C:\JCBroker\lib\jcb.jar` to the classpath option on the Java command line.

Note: It is not recommended that you add it to the system CLASSPATH environment variable, because it might conflict with class loading when running applets that set `applet_jcb.jar` in the ARCHIVE tag.

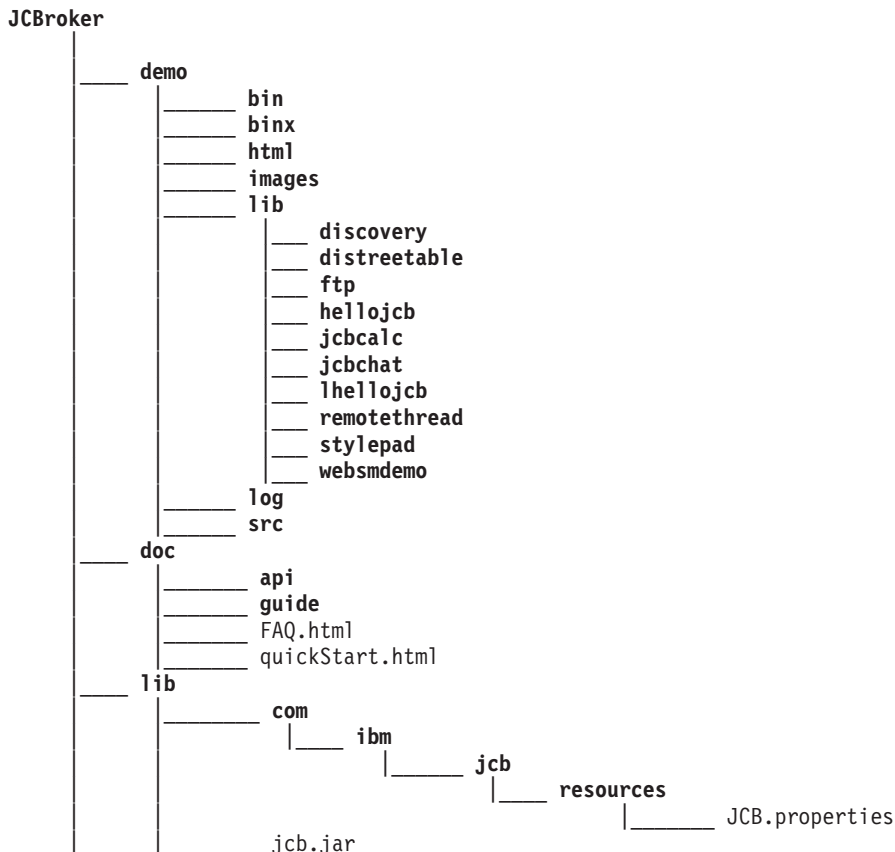
5. To run CBJ in debug mode, replace `jcb.jar` with `jcbd.jar` in the classpath.

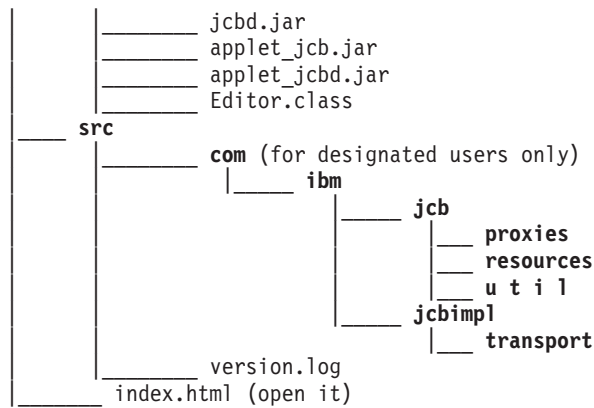
To install CBJ on UNIX^(R), follow the same step for Windows, except for the following system dependent changes:

- Replace the Windows file separator "`\`" with the UNIX file separator "`/`".
 - Replace the Windows classpath separator "`;`" with the UNIX classpath separator "`:`".
 - Replace the Windows system environment variables "`%XXX%`" with the UNIX system environment variables "`$XXX`".
6. For more information on CBJ APIs, running demos, editing properties, designing and writing CBJ applications, and other topics, refer to the `JCBroker/index.html` file in the `cpj_1.1.jar` package.

Package content of `cbj_1.1.jar`

Here is the Class Broker for Java^(TM) (CBJ) package (`cbj_1.1.jar`). The bold items denote main directories of subdirectories. The non-bold items denote files.





The next level of directories under the JCBroker directory consists of the following data:

demo

This directory contains these subdirectories: bin, binx, html, images, lib, log, and src. These directories contain the Windows^(R) .exe files, UNIX^(R) .exe, HTML pages, GIF files, class files, and source code (for some examples) of several demo applications. You can run these demo examples by following the demo instructions. The log directory is a placeholder for debug log files.

doc

This directory contains these two subdirectories: api and guide. The CBJ API guide and a user guide about CBJ are stored here. The API documentation was generated with the J2SDK, javadoc tool. To browse the API document or the user guide, open the index.html file in the corresponding directory. The user guide is an introductory level document, whereas the QuickStart guide is a fast track entry document.

This directory also contains a FAQ.html file and the quickStart.html file that explains how to install and use CBJ, and shows you how to run the demo examples.

lib

This directory contains the CBJ classes and resources that are needed for you to develop, run, and deploy new CBJ-based applications. The classes are packaged in the jcb.jar file, For debug mode, the classes are packaged in the jcbd.jar file. The classes that are included in the ARCHIVE tag of an applet are packaged in the applet_jcb.jar file. For debug mode, the classes are packaged in the applet_jcbd.jar file. The Java Class Broker (JCB) properties file, named JCB.properties, is in the com/ibm/jcb/resources subdirectory. It is read by the CBJ runtime when started. You can change this file to reflect your preferences. The properties are explained in the installation and setting up section. The properties that only apply to applets are contained in the applet_JCB.properties file, which is part of the applet_jcb.jar and applet_jcbd.jar. You can also change this file to reflect your preferences when running in applet mode. To change this file, run the Editor Java application that resides in this directory.

src

This directory contains the CBJ source code and the internal version log file, called version.log. The directory is empty except for selected packages.

index.html

This is the starting page that references you to the rest of the documentation.

Native Abstract Windowing Toolkit

Native Abstract Windowing Toolkit (NAWT) provides Java[®] applications and servlets with the capability to use the Java Development Kit's (JDK) Abstract Windowing Toolkit (AWT) graphics function.

You can use NAWT (pronounced "naw-TEE") in Java graphics applications that do not require direct user interaction. An example of such an application would be code that generates an image such as a JPEG or GIF-encoded file or output stream. For this release, NAWT supports JDK versions 1.2 and 1.3.

You can also use NAWT as an alternative to Remote Abstract Windowing Toolkit (RAWT) for noninteractive applications. In RAWT, graphics computation is performed by an attached PC server. RAWT can be used for a more general class of graphics-based applications, but requires the remote PC, and usually incurs some performance overhead due to network latency between the iSeries server and PC GUI server.

NAWT uses the X Window System as the underlying graphics engine, performing the graphics computation entirely on an iSeries server. The X Window System is a graphics system that provides a client/server base for displaying graphics.


For more information about installing NAWT and using it in your applications, see the following topics:

Install NAWT

Read about installing NAWT on your iSeries server, including information about the recommended X Window graphics server, OS/400 PASE, required software fixes (PTFs), the iSeries Tools for Developers PRPQ, and more.

Tips on using VNC

Find out how to start and VNC display server from a CL program and how to end a VNC display server.

For more information about the VNC server that NAWT uses, refer to the AT&T Research Virtual Network Computing 

Installing NAWT

NAWT relies upon the X Window System for its graphics engine. The recommended X Window graphics server for this release is Virtual Network Computing (VNC), which runs under OS/400 Portable Application Solutions Environment (PASE).

PASE is a UNIX-like environment capable of running most binary executables compiled for IBM's AIX operating system. VNC is a product of AT&T Research Labs, and provides a *virtual* X Window server; that is, an X server that does not require you to directly attach a graphics-capable display device.

To install and run Native Abstract Windowing Toolkit (NAWT), perform the following tasks:

1. Install OS/400 PASE
2. Install NAWT PTFs
3. Install iSeries Tools for Developers PRPQ
4. Create a VNC password file
5. Configure Java system properties
6. Start the VNC server
7. Set environment variables
8. Verify installation procedure

Install OS/400 PASE: Order and install OS/400 Portable Application Solutions Environment (PASE), 5722SS1, option 33. For more information, see OS/400 PASE.

Install NAWT PTFs: Before installing any software fixes (PTFs), ensure that you have the licensed program 5722JV1 option corresponding to the JDK version you want to use by performing these steps:

1. Enter the Go Licensed Program (GO LICPGM) command on the command line.
2. Select option 10 (Displayed installed licensed program) and verify that the licensed program 5722JV1 option corresponding to the version of the JDK you intend to use is installed.


The options are:

JDK Version	Option
JDK 1.2	57SSJV1 option 3
JDK 1.3	57SSJV1 option 5

➤ Apply the latest Java group software fix to pick up any recent NAWT fixes. For more information about software fixes, see Use software fixes.

Install iSeries Tools for Developers PRPQ: Install the iSeries Tools for Developers PRPQ (5799PTL). If you do not have the PRPQ, you must order it.

Newer versions of the PRPQ include a pre-compiled OS/400 PASE-enabled version of Virtual Network Computing (VNC). Older versions do not include VNC. How you install the PRPQ depends on which version you have:

- For versions of the PRPQ ordered on or after 14 June 2002: Complete this task by using the installation instructions available at the Application Factory - iSeries Tools for Development  Web site.

Note: To install the VNC support available in the PRPQ, follow only the installation instructions at the Application Factory Web site. You do not need to follow the setup instructions.

- For versions of the PRPQ ordered before 14 June 2002, refer to Installing older versions of iSeries Tools for Developers PRPQ to complete this task.

Create a VNC password file: By default, VNC requires a password file that protects the VNC display against unauthorized user access. How you create create an encrypted password depends in which version of the PRPQ you are using:

- For versions of the PRPQ ordered on or after 14 June 2002, use the following commands at the iSeries command prompt:

```
MKDIR DIR('/home/your_profile_name/.vnc')
QAPTL/VNCPASSWD USEHOME(*NO) PWDFILE('/home/your_profile_name/.vnc/passwd')
```

- For versions of the PRPQ ordered before 14 June 2002, use the following commands at the iSeries command prompt:

```
MKDIR DIR('/home/your_profile_name/.vnc')
VNCSAVF/VNCPASSWD USEHOME(*NO) PWDFILE('/home/your_profile_name/.vnc/passwd')
```

Only the user that starts the VNC server needs to have a password file. <<

Configure Java system properties

Set the Java system properties. The first line configures Java for the desired JDK version (1.2 or 1.3) and the second line enables NAWT:

```
java.version=version
os400.class.path.rawt=2
```

Where *version* is either 1.2 or 1.3, depending on which version of the JDK you wish to use.

For help on how to set Java system properties, see Customize your iSeries server for the IBM Developer Kit for Java.

Start the VNC server

To start the VNC server, perform the following steps:

1. Start a PASE shell:

```
CALL QP2TERM
```

2. From the PASE shell, start the NAWT-configured version of the VNC server:

```
/QOpenSys/QIBM/ProdData/DeveloperTools/vnc/vncserver_java
```

When the server starts, you will see a message similar to, "New 'X'desktop is *systemname:1*." Remember this desktop name for the next step.

Note: The display number (that is, the number to the right of the colon) may be different if you start another VNC server. Each concurrently running VNC server requires a unique display number. When you do not specify the display number on the call to `vncserver_java`, the `vncserver_java` program finds an available display. Request a specific display by using the following command to start the `vncserver_java` program:

```
/QOpenSys/QIBM/ProdData/DeveloperTools/vnc/vncserver_java :n
```

where *n* is the display number you want to use.

Set environment variables

In any session where you run Java with NAWT, you must tell Java where to find the X server and where to find X authority files. You give Java this information by setting environment variables.

In the session where you want to run Java programs, set the `DISPLAY` environment variable to your system name and display number (that is, the value printed when you ran `vncserver_java`).

Also, set the `XAUTHORITY` environment variable to `/home/your_profile_name/.Xauthority`.

For example, at the iSeries command prompt, type the following commands:

```
ADDENVVAR ENVVAR(DISPLAY) VALUE('systemname:1')
ADDENVVAR ENVVAR(XAUTHORITY) VALUE('/home/your_profile_name/.Xauthority')
```

Notes:

- The `.Xauthority` file is created or modified when the VNC server is started. X server authorization enforces a security protocol on your X server connection, preventing the applications of other users from intercepting your X server requests.
- Only the user who actually starts the JVM needs to set these environment variables. For instance, in a servlet environment, only the user starting the servlet engine needs to set them.

Verify installation procedure

Verify that the steps above completed successfully by using the following command to run a NAWT test Java application:

```
JAVA CLASS(NAWTtest) CLASSPATH('/QIBM/ProdData/Java400/')
```

The test application creates a JPEG-encoded image called `/tmp/NAWTtest.jpg` in the integrated file system. Verify that the application created the file and produced no Java exceptions. If you want to display the image, you can use binary mode to upload the image file to a graphics-capable system.

Installing older versions of iSeries Tools for Developers

➤ Versions of iSeries Tools for Developers PRPQ (5799PTL) ordered before 14 June 2002, PRPQ do not include a pre-compiled OS/400 PASE-enabled version of Virtual Network Computing (VNC).

Use the following instructions to determine if you have the enhanced PRPQ and to install VNC if you have an older version of the PRPQ.

Determine whether you have the enhanced PRPQ: If you own PRPQ 5799-PTL but are not sure whether you have the enhanced version that contains VNC, check for the existence of the following file:


```
/QOpenSys/QIBM/ProdData/DeveloperTools/vnc/vncserver_java
```

The enhanced version of the PRPQ includes the vncserver_java file, but older versions do not. If vncserver_java is not present on your iSeries server, you can either order and install the latest version of the PRPQ or use the following instructions to complete the VNC installation.

Install VNC: To install VNC on an older version of iSeries Tools for Developers PRPQ, complete the following steps.

1. Create the save files on your iSeries server by running the following commands:

```
crtlib vncsavf
crtsavf vncsavf/vncpasswd
crtsavf vncsavf/vnc
crtsavf vncsavf/fonts
crtsavf vncsavf/icewm
```

2. Click the links in the following list to download the save files to your workstation from the Application Factory - iSeries Tools for Development  Web site:

- vnc.savf
- vncpasswd.savf
- fonts.savf
- icewm.savf

3. Use FTP to transfer the save files from your workstation to the iSeries server by running the following commands on your workstation:

```
ftp youriseriesserver
bin
cd /qsys.lib/vncsavf.lib
put vnc.savf
put vncpasswd.savf
put fonts.savf
put icewm.savf
quit
```

4. Restore the save files by running the following commands on your iSeries server:

```
RSTOBJ OBJ(*ALL) SAVLIB(VNCSAVF) DEV(*SAVF) SAVF(VNCSAVF/VNCPASSWD)
RST DEV('/Qsys.lib/vncsavf.lib/vnc.file') OBJ('/QOpenSys/QIBM/ProdData/DeveloperTools/vnc*')
RST DEV('/Qsys.lib/vncsavf.lib/fonts.file') OBJ('/QOpenSys/QIBM/ProdData/DeveloperTools/fonts*')
RST DEV('/Qsys.lib/vncsavf.lib/icewm.file') OBJ('/QOpenSys/QIBM/ProdData/DeveloperTools/icewm*')
```

5. Continue installing NAWT by creating a VNC password file. 

Tips on using VNC

 This page discusses additional tips on using Virtual Network Computing (VNC).

Start a VNC display server from a CL program

The following example is one way to set the DISPLAY environment variable and start VNC automatically from a CL program. It is assumed that display :n is not already running. It is also assumed that you have already created the password file required by VNC by running the VNCPASSWD command.

```
ADDENVVAR ENVVAR(DISPLAY) VALUE('systemname:n')
call qp2shell parm('/QOpenSys/QIBM/ProdData/DeveloperTools/vnc/vncserver_java' ':n')
```

where *n* is the numeric value representing the display number that you want to start.

End a VNC display server

To end a VNC server that you started:

```
call qp2shell parm('/QOpenSys/QIBM/ProdData/DeveloperTools/vnc/vncserver_java' '-kill' ':n')
```

where n is the numeric value representing the display number that you want to end. <<

Chapter 2. Java security

The majority of the Java^(TM) programs that run on an iSeries server are applications, not applets, so the "sandbox" security model does not restrict them. From a security point of view, Java applications are subject to the same security restrictions as any other program on an iSeries server. To run a Java program on an iSeries server, you must have authority to the class file in the integrated file system. Once the program starts, it runs under the user's authority.

You can use adopted authority to access objects with the authority of the user that is running the program, and the program owner's authority. Adopted authority temporarily gives a user authority to objects that they would not have originally had authority to access. See the Create Java Program (CRTJVAPGM) command information for details on the two new adopted authority parameters, which are USRPRF and USEADPAUT.

➤ The IBM Developer Kit for Java provides the following security features for your Java applications:

Java security model

The byte code loader and verifier, within the Java virtual machine, also provide a measure of Java security using the Java security model. Just as with applets, the byte code loader and verifier check that the byte codes are valid and data types are used properly. They also check that registers and memory are accessed correctly, and that the stack does not overflow or underflow. These checks ensure that the Java virtual machine can safely run the class without compromising the integrity of the system.

Java Cryptography Extension

The Java Cryptography Extension (JCE) implementation on the iSeries server is compatible with the implementation of Sun Microsystems, Inc. This documentation covers the unique aspects of the iSeries implementation. It is assumed that you are familiar with the general documentation for the JCE.

Java Secure Socket Extension

The Java Secure Socket Extension (JSSE) is the Java implementation of the Secure Sockets Layer (SSL) protocol. JSSE uses SSL and the Transport Layer Security (TLS) protocol to enable clients and servers to conduct secure communications over TCP/IP. This documentation covers the unique aspects of the iSeries implementation of JSSE. It is assumed that you are familiar with the general documentation for JSSE.

Java Authentication and Authorization Service

Java Authentication and Authorization Service (JAAS) is another item of security that the IBM Developer Kit for Java supports. Currently, the Java 2 Software Development Kit (J2SDK), Standard Edition provides access controls that are based on where the code originated and who signed the code (code source-based access controls). It lacks, however, the ability to enforce additional access controls based on who runs the code. JAAS provides a framework that adds this support to the Java 2 security model.

Java Generic Security Service

Java Generic Security Service (JGSS) is another item of security that the IBM Developer Kit for Java supports. JGSS provides a generic interface for secure messaging between applications. JGSS supports a variety of security mechanisms based on secret-key, public-key, or other security technologies.

Note: For J2SDK, version 1.4, JAAS, JCE, JGSS, and JSSE are part of the base JDK and are not considered to be extensions. For previous JDK versions, these security items are extensions. ⏪

Java security model

You can download Java^(TM) applets from any system; thus, security mechanisms exist within the Java virtual machine to protect against malicious applets. The Java runtime system verifies the bytecodes as the Java virtual machine loads them. This ensures that they are valid bytecodes and that the code does not violate any of the restrictions that the Java virtual machine places on Java applets. Java applets are restricted in what operations they can perform, how they access memory, and how they use the Java virtual machine. The restrictions are in place to prevent a Java applet from gaining access to underlying operating system or data on the system. This is the "sandbox" security model, because the Java applet can only "play" in its own sandbox.

The "sandbox" security model is a combination of the class loader, class file verifier, and the `java.lang.SecurityManager` class.

For more information about security, see the Security by Sun Microsystems, Inc. documentation and Secure applications with SSL.

Java Cryptography Extension

The Java^(TM) Cryptography Extension (JCE) 1.2 is a standard extension to the Java 2 Software Development Kit (J2SDK), Standard Edition. The JCE implementation on an iSeries server is compatible with the implementation of Sun Microsystems, Inc. This documentation covers the unique aspects of the iSeries implementation. We assume that you are familiar with the general documentation for the JCE extensions. To make it easier for you to work with that and our iSeries information, we provide a link to the Sun JCE documentation [👉](#).

On the iSeries server, the level of encryption is controlled by the Cryptographic Access Provider Product. This is available in two versions, 5722-AC2 and 5722-AC3. The 5722-AC3 product allows all of the encryption algorithms. The 5722-AC2 product does not allow Triple-DES, and limits symmetric algorithms to 56 bits and asymmetric algorithms to 1024 bits.

With the exception of the restrictions on 5722-AC2 already mentioned, the IBM JCE Provider supports the following algorithms:

- DES
- Triple-DES
- RC2
- RC4
- Blowfish
- RSA
- Diffie-Hellman
- DSA
- Mars
- MD2
- MD5
- SHA-1
- Seal

In addition, it also provides a random number generator.

➤ If you want to use IBM JCE with Java 1.2, edit the `/QIBM/ProdData/Java400/jdk12/lib/security/java.security` file. ⏪ The section of the file that needs to be changed is shown as follows.

```
#
# To use the IBMJCE security provider, you need to:
# 1) Install an IBM Cryptographic Access Provider Product
# 2) uncomment the second provider entry that follows.
#
# List of providers and their preference orders:
#
security.provider.1=sun.security.provider.Sun
#security.provider.2=com.ibm.crypto.provider.IBMJCE
```

➤ If you want to use IBM JCE with Java 1.3, edit the /QIBM/ProdData/OS400/Java400/jdk/lib/security/java.security file. ⚡ The section of the file that needs to be changed is shown as follows.

```
#
# To use the IBMJCE security provider, you need to:
# 1) Install an IBM Cryptographic Access Provider Product
# 2) Uncomment the third provider entry that follows.
#
# List of providers and their preference orders:
#
security.provider.1=sun.security.provider.Sun
security.provider.2=com.sun.rsa.jca.Provider
#security.provider.3=com.ibm.crypto.provider.IBMJCE
```

In both cases, it is just a matter of deleting one character.

Java Secure Socket Extension

The Java[™] Secure Socket Extension (JSSE) is the Java implementation of the Secure Sockets Layer (SSL) protocol. JSSE uses SSL and the Transport Layer Security (TLS) protocol to enable clients and servers to conduct secure communications over TCP/IP.

JSSE provides the following functions:

- Encrypts data
- Authenticates remote user IDs
- Authenticates remote system names
- Performs client/server authentication
- Ensures message integrity



Integrated into the Java 2 Software Development Kit, Standard Edition (J2SDK), version 1.4, JSSE provides more functionality than does SSL alone. For more information, see the following topics:

Using SSL (JSSE, version 1.0.8)

SSL provides a means of authenticating a server and a client to provide privacy and data integrity. All SSL communications begin with a "handshake" between the server and the client. During the handshake, SSL negotiates the cipher suite that the client and server use to communicate with each other. This cipher suite is a combination of the various security features available through SSL. You can use SSL with any supported version of the JDK prior to J2SDK, version 1.4.

Using JSSE, version 1.4

JSSE is like a framework that abstracts the underlying mechanisms of both SSL and TLS. By abstracting the complexity and peculiarities of the underlying protocols, JSSE enables programmers to use secure, encrypted communications while at the same time minimizing possible security vulnerabilities. This information applies only to using JSSE on iSeries servers that run J2SDK, version 1.4.

Note: This information concerns the version of JSSE that now comes bundled in the J2SDK, version 1.4. For previous versions of JSSE, see Java Secure Socket Extension on the Sun Java Web site  . 

Using SSL (JSSE, version 1.0.8)

You can use the Java Secure Socket Extension (JSSE, version 1.0.8), which is the Java implementation of secure sockets layer (SSL), to make your Java[™] application more secure. SSL does the following to improve the security of your application:




- Protects communication data through encryption.
- Authenticates remote user IDs.
- Authenticates remote system names.

Note: SSL uses a digital certificate to encrypt the socket communication of your Java application. Digital certificates are an Internet standard for identifying secure systems, users, and applications. You can control digital certificates using the IBM Digital Certificate Manager. For more information, see IBM Digital Certificate Manager.

To make your Java application more secure by using SSL:

- Prepare the iSeries server to support SSL.
- Design your Java application to use SSL by:
 - Changing your Java socket code to use socket factories if you do not use socket factories already.
 - Changing your Java code to use SSL.
- Use a digital certificate to make your Java application more secure by:
 1. Selecting a type of digital certificate to use.
 2. Using the digital certificate when you run your application.

You can also register your Java application as a secure application by using the QsyRegisterAppForCertUse API. For more information, see QsyRegisterAppForCertUse.



 For more information on the Java version of SSL, see Java Secure Socket Extension  by Sun Microsystems, Inc. 

Prepare iSeries server for secure sockets layer support

To prepare your system to use secure sockets layer (SSL), you need to install the Digital Certificate Manager LP:

-  5722-SS1 OS/400 - Digital Certificate Manager 

You also need to install one of these Cryptographic Access Provider LPs:

-  5722-AC1 Cryptographic Access Provider 40-Bit
- 5722-AC2 Cryptographic Access Provider 56-Bit
- 5722-AC3 Cryptographic Access Provider 128-Bit 

You also need to make sure you can access or create a digital certificate on your system. For more information on iSeries digital certificate management and the Internet, see Getting started with IBM Digital Certificate Manager.

Cryptographic Access Providers

The Cryptographic Access Providers offer many cipher suites for your system to use. A cipher suite is a combination of different security features. This list shows which cipher suite each Cryptographic Access Provider offers:

» 5722-AC1 Cryptographic Access Provider 40-Bit «

```
SSL_RSA_WITH_NULL_MD5
SSL_RSA_WITH_NULL_SHA
SSL_RSA_EXPORT_WITH_RC4_40_MD5
SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5
```

» 5722-AC2 Cryptographic Access Provider 56-Bit «

```
SSL_RSA_WITH_NULL_MD5
SSL_RSA_WITH_NULL_SHA
SSL_RSA_WITH_DES_CBC_SHA
SSL_RSA_WITH_DES_CBC_MD5
SSL_RSA_EXPORT_WITH_RC4_40_MD5
SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5
```

» 5722-AC3 Cryptographic Access Provider 128-Bit «

```
SSL_RSA_WITH_NULL_MD5
SSL_RSA_WITH_NULL_SHA
SSL_RSA_EXPORT_WITH_RC4_40_MD5
SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5
SSL_RSA_WITH_RC4_128_SHA
SSL_RSA_WITH_DES_CBC_SHA
SSL_RSA_WITH_3DES_EDE_CBC_SHA
SSL_RSA_WITH_RC4_128_MD5
SSL_RSA_WITH_RC2_CBC_128_MD5
SSL_RSA_WITH_DES_CBC_MD5
SSL_RSA_WITH_3DES_EDE_CBC_MD5
```

You may be limited on the Cryptographic Access Provider that you can choose from, depending on which country or region you are in. Once you load a Cryptographic Access Provider, you can use any of the cipher suites that that Access Provider offers.

Change your Java code to use socket factories

To use secure sockets layer (SSL) with your existing code, you must first change your code to use socket factories.

To change your code to use socket factories, perform the following steps:

1. Add this line to your program to import the `SocketFactory` class:

```
import javax.net.*;
```

2. Add a line that declares an instance of a `SocketFactory` object. For example:

```
SocketFactory socketFactory
```

3. Initialize the `SocketFactory` instance by setting it equal to the method `SocketFactory.getDefault()`. For example:

```
socketFactory = SocketFactory.getDefault();
```

The whole declaration of the `SocketFactory` should look like this:

```
SocketFactory socketFactory = SocketFactory.getDefault();
```

4. Initialize your existing sockets. Call the `SocketFactory` method `createSocket(host,port)` on your socket factory for each socket you declare.

Your socket declarations should now look like this:

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

Where:

- `s` is the socket that is being created.
- `socketFactory` is the `SocketFactory` that was created in step 2.

- *host* is a string variable that represents the name of a host server.
- *port* is an integer variable that represents the port number of the socket connection.

When you have completed all of these steps, your code uses socket factories. You do not need to make any other changes to your code. All of the methods that you call and all the syntax with your sockets still work.

See Examples: Change your Java^(TM) code to use server socket factories for an example of a client program being converted to use socket factories.

See Example: Change your Java code to use client socket factories for an example of a client program being converted to use socket factories.

Examples: Change your Java code to use server socket factories

These examples show you how to change a simple socket class, named `simpleSocketServer`, so that it uses socket factories to create all of the sockets. The first example shows you the `simpleSocketServer` class without socket factories. The second example shows you the `simpleSocketServer` class with socket factories. In the second example, `simpleSocketServer` is renamed to `factorySocketServer`.

Example 1: Socket server program without socket factories

Note: Read the Code example disclaimer for important legal information.

```

/* File simpleSocketServer.java*/

import java.net.*;
import java.io.*;

public class simpleSocketServer {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java simpleSocketServer serverPort");
            System.out.println("Defaulting to port 3000 since serverPort not specified.");
        }
        else
            serverPort = new Integer(args[0]).intValue();

        System.out.println("Establishing server socket at port " + serverPort);

        ServerSocket    serverSocket =
            new ServerSocket(serverPort);

        // a real server would handle more than just one client like this...

        Socket s = serverSocket.accept();
        BufferedInputStream is = new BufferedInputStream(s.getInputStream());
        BufferedOutputStream os = new BufferedOutputStream(s.getOutputStream());

        // This server just echoes back what you send it...

        byte buffer[] = new byte[4096];

        int bytesRead;

        // read until "eof" returned
        while ((bytesRead = is.read(buffer)) > 0) {
            os.write(buffer, 0, bytesRead); // write it back
            os.flush(); // flush the output buffer
        }
    }
}

```

```

        s.close();
        serverSocket.close();
    } // end main()
} // end class definition

```

Example 2: Simple socket server program with socket factories

Note: Read the Code example disclaimer for important legal information.

```

/* File factorySocketServer.java */

// need to import javax.net to pick up the ServerSocketFactory class
import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySocketServer {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java simpleSocketServer serverPort");
            System.out.println("Defaulting to port 3000 since serverPort not specified.");
        }
        else
            serverPort = new Integer(args[0]).intValue();

        System.out.println("Establishing server socket at port " + serverPort);

        // Change the original simpleSocketServer to use a
        // ServerSocketFactory to create server sockets.
        ServerSocketFactory serverSocketFactory =
            ServerSocketFactory.getDefault();
        // Now have the factory create the server socket. This is the last
        // change from the original program.
        ServerSocket serverSocket =
            serverSocketFactory.createServerSocket(serverPort);

        // a real server would handle more than just one client like this...

        Socket s = serverSocket.accept();
        BufferedInputStream is = new BufferedInputStream(s.getInputStream());
        BufferedOutputStream os = new BufferedOutputStream(s.getOutputStream());

        // This server just echoes back what you send it...

        byte buffer[] = new byte[4096];

        int bytesRead;

        while ((bytesRead = is.read(buffer)) > 0) {
            os.write(buffer, 0, bytesRead);
            os.flush();
        }

        s.close();
        serverSocket.close();
    }
}

```

For background information, see [Change your Java^{\(TM\)} code to use socket factories](#).

Examples: Change your Java code to use client socket factories

These examples show you how to change a simple socket class, named `simpleSocketClient`, so that it uses socket factories to create all of the sockets. The first example shows you the `simpleSocketClient` class without socket factories. The second example shows you the `simpleSocketClient` class with socket factories. In the second example, `simpleSocketClient` is renamed to `factorySocketClient`.

Example 1: Socket client program without socket factories

Note: Read the Code example disclaimer for important legal information.

```
/* Simple Socket Client Program */

import java.net.*;
import java.io.*;

public class simpleSocketClient {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java simpleSocketClient serverHost serverPort");
            System.out.println("serverPort defaults to 3000 if not specified.");
            return;
        }
        if (args.length == 2)
            serverPort = new Integer(args[1]).intValue();

        System.out.println("Connecting to host " + args[0] + " at port " +
            serverPort);

        // Create the socket and connect to the server.
        Socket s = new Socket(args[0], serverPort);
        .
        .
        .

        // The rest of the program continues on from here.
    }
}
```

Example 2: Simple socket client program with socket factories

Note: Read the Code example disclaimer for important legal information.

```
/* Simple Socket Factory Client Program */

// Notice that javax.net.* is imported to pick up the SocketFactory class.
import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySocketClient {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java factorySocketClient serverHost serverPort");
            System.out.println("serverPort defaults to 3000 if not specified.");
            return;
        }
        if (args.length == 2)
            serverPort = new Integer(args[1]).intValue();

        System.out.println("Connecting to host " + args[0] + " at port " +
            serverPort);
    }
}
```

```

// Change the original simpleSocketClient program to create a
// SocketFactory and then use the socket factory to create sockets.

SocketFactory socketFactory = SocketFactory.getDefault();

// Now the factory creates the socket. This is the last change
// to the original simpleSocketClient program.

Socket s = socketFactory.createSocket(args[0], serverPort);
.
.
.

// The rest of the program continues on from here.

```

For background information, see [Change your Java^{\(TM\)} code to use socket factories](#).

Change your Java code to use secure sockets layer

If your code already uses socket factories to create its sockets, then you can add secure socket layer (SSL) support to your program. If your code does not already use socket factories, see [Change your Java^{\(TM\)} code to use socket factories](#).

To change your code to use SSL, perform the following steps:

1. Import `javax.net.ssl.*` to add SSL support:

```
import javax.net.ssl.*;
```

2. Declare a `SocketFactory` by using `SSLSocketFactory` to initialize it:

```
SocketFactory newSF = SSLSocketFactory.getDefault();
```

3. Use your new `SocketFactory` to initialize your sockets the same way that you used your old `SocketFactory`:

```
Socket s = newSF.createSocket(args[0], serverPort);
```

Your code now uses SSL support. You do not need to make any other changes to your code.

See [Examples: Change your Java client to use secure sockets layer](#) and [Examples: Change your Java server to use secure sockets layer](#) for example code.

Examples: Change your Java server to use secure sockets layer

These examples show you how to change one class, named `factorySocketServer`, to use secure sockets layer (SSL).

The first example shows you the `factorySocketServer` class not using SSL. The second example shows you the same class, renamed `factorySSLSocketServer`, using SSL.

Example 1: Simple `factorySocketServer` class without SSL support

Note: Read the [Code example disclaimer](#) for important legal information.

```

/* File factorySocketServer.java */
// need to import javax.net to pick up the ServerSocketFactory class
import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySocketServer {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java simpleSocketServer serverPort");

```

```

    System.out.println("Defaulting to port 3000 since serverPort not specified.");
}
else
    serverPort = new Integer(args[0]).intValue();

System.out.println("Establishing server socket at port " + serverPort);

// Change the original simpleSocketServer to use a
// ServerSocketFactory to create server sockets.
ServerSocketFactory serverSocketFactory =
    ServerSocketFactory.getDefault();
// Now have the factory create the server socket. This is the last
// change from the original program.
ServerSocket serverSocket =
    serverSocketFactory.createServerSocket(serverPort);

// a real server would handle more than just one client like this...

Socket s = serverSocket.accept();
BufferedInputStream is = new BufferedInputStream(s.getInputStream());
BufferedOutputStream os = new BufferedOutputStream(s.getOutputStream());

// This server just echoes back what you send it.

byte buffer[] = new byte[4096];

int bytesRead;

while ((bytesRead = is.read(buffer)) > 0) {
    os.write(buffer, 0, bytesRead);
    os.flush();
}

s.close();
serverSocket.close();
}
}

```

Example 2: Simple factorySocketServer class with SSL support

Note: Read the Code example disclaimer for important legal information.

```

/* File factorySocketServer.java */

// need to import javax.net to pick up the ServerSocketFactory class
import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySocketServer {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java simpleSocketServer serverPort");
            System.out.println("Defaulting to port 3000 since serverPort not specified.");
        }
        else
            serverPort = new Integer(args[0]).intValue();

        System.out.println("Establishing server socket at port " + serverPort);

        // Change the original simpleSocketServer to use a
        // ServerSocketFactory to create server sockets.
        ServerSocketFactory serverSocketFactory =
            ServerSocketFactory.getDefault();

```

```

// Now have the factory create the server socket. This is the last
// change from the original program.
ServerSocket serverSocket =
    serverSocketFactory.createServerSocket(serverPort);

// a real server would handle more than just one client like this...

Socket s = serverSocket.accept();
BufferedInputStream is = new BufferedInputStream(s.getInputStream());
BufferedOutputStream os = new BufferedOutputStream(s.getOutputStream());

// This server just echoes back what you send it.

byte buffer[] = new byte[4096];

int bytesRead;

while ((bytesRead = is.read(buffer)) > 0) {
    os.write(buffer, 0, bytesRead);
    os.flush();
}

s.close();
serverSocket.close();
}
}

```

For background information, see [Change your JavaTM code to use secure sockets layer](#).

Examples: Change your Java client to use secure sockets layer

These examples show you how to change one class, named `factorySocketClient`, to use secure sockets layer (SSL).

The first example shows you the `factorySocketClient` class not using SSL. The second example shows you the same class, renamed `factorySSLSocketClient`, using SSL.

Example 1: Simple `factorySocketClient` class without SSL support

Note: Read the [Code example disclaimer](#) for important legal information.

```

/* Simple Socket Factory Client Program */

import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySocketClient {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java factorySocketClient serverHost serverPort");
            System.out.println("serverPort defaults to 3000 if not specified.");
            return;
        }
        if (args.length == 2)
            serverPort = new Integer(args[1]).intValue();

        System.out.println("Connecting to host " + args[0] + " at port " +
            serverPort);

        SocketFactory socketFactory = SocketFactory.getDefault();

        Socket s = socketFactory.createSocket(args[0], serverPort);
    }
}

```

```

.
:
.

// The rest of the program continues on from here.

```

Example 2: Simple factorySocketClient class with SSL support

Note: Read the Code example disclaimer for important legal information.

```

// Notice that we import javax.net.ssl.* to pick up SSL support
import javax.net.ssl.*;
import javax.net.*;
import java.net.*;
import java.io.*;

public class factorySSLSocketClient {
    public static void main (String args[]) throws IOException {

        int serverPort = 3000;

        if (args.length < 1) {
            System.out.println("java factorySSLSocketClient serverHost serverPort");
            System.out.println("serverPort defaults to 3000 if not specified.");
            return;
        }
        if (args.length == 2)
            serverPort = new Integer(args[1]).intValue();

        System.out.println("Connecting to host " + args[0] + " at port " +
            serverPort);

        // Change this to create an SSLSocketFactory instead of a SocketFactory.
        SocketFactory socketFactory = SSLSocketFactory.getDefault();

        // We do not need to change anything else.
        // That's the beauty of using factories!
        Socket s = socketFactory.createSocket(args[0], serverPort);
        .
        .
        .

        // The rest of the program continues on from here.

```

For background information, see [Change your Java^{\(TM\)} code to use secure sockets layer](#).

Select a digital certificate to use

You should consider several factors when deciding which digital certificate to use. You can use your system's default certificate or you can specify another certificate to use.

You want to use your system's default certificate if:

- You do not have any specific security requirements for your Java^(TM) application.
- You do not know what kind of security you need for your Java application.
- Your system's default certificate meets the security requirements for your Java application.

Note: If you decide that you want to use your system's default certificate, check with your system administrator to make sure that a default system certificate has been created. For more information on digital certificate management, see [Getting started with IBM Digital Certificate Manager](#).

If you do not want to use your system's default certificate, you need to choose a different certificate to use. You can choose from two types of certificates:

- **User certificate** that identifies the user of the application.

- **System certificate** that identifies the system on which the application is running.

➤ You should use a user certificate if:

- your application runs as a client application.
- you want the certificate to identify the user who is working with the application.

You should use a system certificate if:

- your application runs as a server application.
- you want the certificate to identify on which system the application is running. ◀◀

Once you know what kind of certificate you need, you can choose from any of the digital certificates in any of the certificate containers that you are able to access.

Use the digital certificate when you run your Java application

To use secure sockets layer (SSL), you must run your Java application using a digital certificate.

To specify which digital certificate to use, use the following properties:

- `os400.certificateContainer`
- `os400.certificateLabel`

For example, if you want run the Java application `MyClass.class` using the digital certificate `MYCERTIFICATE`, and `MYCERTIFICATE` was in the digital certificate container `YOURDCC`, then the `java` command would look like this:

```
java -Dos400.certificateContainer=YOURDCC
     -Dos400.certificateLabel=MYCERTIFICATE MyClass
```

If you have not already decided which digital certificate to use, see [Select a digital certificate to use](#). You may also decide to use your system's default certificate, which is stored in the system's default certificate container.

To use your system's default digital certificate, you do not need to specify a certificate or a certificate container anywhere. Your Java application uses your system's default digital certificate automatically.

For more information on iSeries digital certificate management and the Internet, see [Getting started with IBM Digital Certificate Manager](#).

Digital certificates and the `-os400.certificateLabel` property: Digital certificates are an Internet standard for identifying secure systems, users, and applications. Digital certificates are stored in digital certificate containers. If you want to use a digital certificate container's default certificate, you do not need to specify a certificate label. If you want to use a specific digital certificate, you must specify that certificate's label in the `java` command using this property:

```
os400.certificateLabel=
```

For example, if the name of the certificate you want to use is `MYCERTIFICATE`, then the `java` command you enter would look like this:

```
java -Dos400.certificateLabel=MYCERTIFICATE MyClass
```

In this example, the Java application `MyClass` would use the certificate `MYCERTIFICATE`. `MYCERTIFICATE` would need to be in the system's default certificate container to be used by `MyClass`.

Digital certificate containers and the `-os400.certificateContainer` property: Digital certificate containers store digital certificates. If you want to use the iSeries system default certificate container, you do not need to specify a certificate container. To use a specific digital certificate container, you need to specify that digital certificate container in the `java` command using this property:

```
os400.certificateContainer=
```

For example, if the name of the certificate container that contains the digital certificate you want to use is named MYDCC, then the `java` command you enter would look like this:

```
java -Dos400.certificateContainer=MYDCC MyClass
```

In this example, the Java application, named `MyClass.class`, would run on the system by using the default digital certificate that is in the digital certificate container named MYDCC. Any sockets that you create in the application use the default certificate that is in MYDCC to identify themselves and make all of their communications secure.

If you wanted to use the digital certificate MYCERTIFICATE in the digital certificate container, then the `java` command that you would enter would look like this:

```
java -Dos400.certificateContainer=MYDCC  
-Dos400.certificateLabel=MYCERTIFICATE MyClass
```



Using Java Secure Socket Extension, version 1.4

Java Secure Socket Extension (JSSE) uses both the Secure Sockets Layer (SSL) protocol, and the Transport Layer Security (TLS) protocol to provide secure, encrypted communications between your clients and servers.

The IBM implementation of JSSE is called IBM JSSE. IBM JSSE includes a native iSeries JSSE provider and a pure Java JSSE provider.

For more information about configuring your iSeries server to support JSSE, use the following links:


Configuring your server to support JSSE

Find out how to configure your iSeries server to use IBM JSSE. Information includes software requirements, how to change JSSE providers, and the necessary security properties and system properties.

Using the native iSeries JSSE provider

Read about using the native iSeries implementations of the JSSE `KeyStore` class and the `SSLConfiguration` class.

JSSE examples

Use the example programs to discover how you can use JSSE in your applications. The example Java source code shows how clients and servers can use `SSLContext` objects on both clients and servers to create a secure communication environment. 






Configuring your iSeries server to support JSSE

When you use the Java 2 Software Development Kit (J2SDK), version 1.4 on your iSeries server, JSSE is already configured. The default configuration uses the native iSeries JSSE provider.

Software requirements: To use JSSE with J2SDK, version 1.4, you must have installed IBM Cryptographic Access Provider 128-bit (5722-AC3) on your iSeries server. For more information, see [Cryptographic Access Providers](#).

Changing JSSE providers: You can configure JSSE to use the pure Java JSSE provider instead of the native iSeries JSSE provider. By changing some specific JSSE security properties and Java system properties, you can switch between the two providers. For more information, see the following topics:

- JSSE providers
- JSSE security properties
- Java system properties

Security managers: If you are running your JSSE application with a Java security manager enabled, you may need to set the available network permissions. For more information, see [SSLPermission](#) in [Permissions in the Java 2 SDK](#)   

JSSE providers

IBM JSSE includes a native iSeries JSSE provider and a pure Java JSSE provider. The provider that you choose to use depends on the needs of your application.

Both JSSE providers adhere to the JSSE interface specification. The two providers can communicate with each other and with any other SSL or TLS implementation (even non-Java implementations).

Pure Java JSSE provider: The pure Java JSSE provider offers the following features:

- Works with any type of KeyStore object to control and configure digital certificates (for example, JKS, PKCS12, and so on)
- Allows you to use any combination of JSSE components from multiple implementations together (for example, you can use an X509TrustManager from the native iSeries JSSE provider to initialize an SSLContext from this pure Java JSSE provider).

IBMJSSE is the provider name for the pure Java implementation. You need to pass this provider name, using the proper case, to the `java.security.Security.getProvider()` method or the various `getInstance()` methods for several of the JSSE classes.

Native iSeries JSSE provider: The native iSeries JSSE provider offers the following features:

- Uses the native iSeries SSL support
- Uses Digital Certificate Manager to configure and control digital certificates
- Offers best performance


Note: The native iSeries JSSE provider requires a unique iSeries type of KeyStore. The JSSE native iSeries JSSE provider also does not allow components from any other implementation to be plugged in to its implementation.

`IbmIseriesSslProvider` is the name for the native iSeries implementation. You need to pass this provider name, using the proper case, to the `java.security.Security.getProvider()` method or the various `getInstance()` methods for several of the JSSE classes.

Changing the default JSSE provider: You can change the default JSSE provider by making the appropriate changes to your security properties. For more information, see the following topic:

- JSSE security properties

After changing the JSSE provider, ensure that your system properties specify the proper configuration for digital certificate information (keystore) required by the new provider. For more information, see the following topic:

- Java system properties 



JSSE security properties

A Java virtual machine (JVM) uses many important security properties that you set by editing the Java master security properties file. This file, named `java.security`, usually resides in the `/QIBM/ProdData/Java400/jdk14/lib/security` directory on your iSeries server.

The following list describes several relevant security properties for using JSSE. Use the descriptions as a guide for editing the `java.security` file.

security.provider.<integer>

The JSSE provider that you want to use. Also statically registers cryptographic provider classes. Specify the different JSSE providers exactly like the following example:

```
security.provider.5=com.ibm.as400.ibmonly.net.ssl.Provider
security.provider.6=com.ibm.jsse.IBMJSSEProvider
```

ssl.KeyManagerFactory.algorithm

Specifies the default `KeyManagerFactory` algorithm. For the native iSeries JSSE provider, use the following:

```
ssl.KeyManagerFactory.algorithm=IbmISeriesX509
```

For the pure Java JSSE provider, use the following:

```
ssl.KeyManagerFactory.algorithm=IbmX509
```

For more information, see the javadoc for `javax.net.ssl.KeyManagerFactory`.

ssl.TrustManagerFactory.algorithm

Specifies the default `TrustManagerFactory` algorithm. For the native iSeries JSSE provider, use the following:

```
ssl.TrustManagerFactory.algorithm=IbmISeriesX509
```

For the pure Java JSSE provider, use the following:

```
ssl.TrustManagerFactory.algorithm=IbmX509
```

For more information, see the javadoc for `javax.net.ssl.TrustManagerFactory`.

ssl.SocketFactory.provider

Specifies the default SSL socket factory. For the native iSeries JSSE provider, use the following:

```
ssl.SocketFactory.provider=com.ibm.as400.ibmonly.net.ssl.SSLSocketFactoryImpl
```

For the pure Java JSSE provider, use the following:

```
ssl.SocketFactory.provider=com.ibm.jsse.JSSESocketFactory
```

For more information, see the javadoc for `javax.net.ssl.SSLSocketFactory`.

ssl.ServerSocketFactory.provider

Specifies the default SSL server socket factory. For the native iSeries JSSE provider, use the following:

```
ssl.ServerSocketFactory.provider=com.ibm.as400.ibmonly.net.ssl.SSLServerSocketFactoryImpl
```

For the pure Java JSSE provider, use the following:

```
ssl.ServerSocketFactory.provider=com.ibm.jsse.JSSEServerSocketFactory
```

For more information, see the javadoc for `javax.net.ssl.SSLServerSocketFactory`.



JSSE Java system properties

To use JSSE in your applications, you need to specify several system properties that the default SSLContext objects needs in order to provide confirmation of the configuration. Some of the properties apply to both providers, while others apply to only the native iSeries provider.

When using the native iSeries JSSE provider, when you specify none of the properties, the `os400.certificateContainer` defaults to `*SYSTEM`, which means that JSSE uses the default entry in the system certificate store.

Properties that work for both providers: The following properties apply to both JSSE providers. Each description includes the default property, if applicable.

javax.net.ssl.trustStore

The name of the file that contains the KeyStore object that you want the default TrustManager to use. The default value is `jssecacerts`, or `cacerts` (if `jssecacerts` does not exist).

javax.net.ssl.trustStoreType

The type of KeyStore object that you want the default TrustManager to use. The default value is the value returned by the `KeyStore.getDefaultType` method.

javax.net.ssl.trustStorePassword

The password for the KeyStore object that you want the default TrustManager to use.

javax.net.ssl.keyStore

The name of the file that contains the KeyStore object that you want the default KeyManager to use.

javax.net.ssl.keyStoreType

The type of KeyStore object that you want the default KeyManager to use. The default value is the value returned by the `KeyStore.getDefaultType` method.

javax.net.ssl.keyStorePassword

The password for the KeyStore object that you want the default KeyManager to use.

Properties that work for the iSeries native JSSE provider only: The following properties apply to the native iSeries JSSE provider only.

os400.secureApplication

The application identifier. JSSE uses this property only when you do not specify any of the following properties:

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

os400.certificateContainer

The name of the keyring that you want to use. JSSE uses this property only when you do not specify any of the following properties:


- `javax.net.ssl.keyStore`

- javax.net.ssl.keyStoreType
- javax.net.ssl.keyStorePassword
- os400.secureApplication

os400.certificateLabel

The keyring label that you want to use. JSSE uses this property only when you set and use the os400.certificateContainer property.

Additional information: For more information about system properties, see the following topics:

- Java system properties for J2SDK, version 1.4, on iSeries servers
- System Properties on the Sun Java Web site 



Using the native iSeries JSSE provider

The native iSeries JSSE provider offers the full suite of JSSE classes and interfaces. To use the native iSeries provider effectively, please refer to the following information:

- Protocol values for the SSLContext.getInstance method
- The native KeyStore implementation
- Restrictions when using the native iSeries provider
- Javadoc information for SSLConfiguration

Protocol values for the SSLContext.getInstance method: The following table identifies and describes the protocol values for the SSLContext.getInstance method of the native iSeries JSSE provider.

Protocol value	Supported SSL protocols
SSL	SSL version 2, SSL version 3, and TLS version 1
SSLv2	SSL version 2
SSLv3	SSL version 3
TLS	SSL version 2, SSL version 3, and TLS version 1
TLSv1	TLS version 1
SSL_TLS	SSL version 2, SSL version 3, and TLS version 1

Native iSeries KeyStore implementation: The native iSeries provider offers an implementation of the KeyStore class of type IbmIseriesKeyStore. This keystore implementation provides a wrapper around the Digital Certificate Manager support. The contents of the keystore are based on a particular application identifier or keyring file, password, and label. JSSE loads the keystore entries from the Digital Certificate Manager. To load the entries, JSSE uses the appropriate application identifier or keyring information when your application makes the first attempt to access keystore entries or keystore information. You can not modify the keystore, and you must make all configuration changes by using must be made by using the Digital Certificate Manager.

For more information about using the Digital Certificate Manager, see the following topic:

Digital Certificate Manager

Restrictions when using the native iSeries provider: For the native iSeries JSSE provider to work, your JSSE application must use only components from the native implementation. For example, your native iSeries JSSE-enabled application can not use an X509KeyManager object created by using the pure Java JSSE provider to successfully initialize an SSLContext object created by using the native iSeries JSSE provider.

Additionally, you have to initialize the implementations of X509KeyManager and X509TrustManager in the native iSeries provider by using either an IbmIseriesKeyStore object or a com.ibm.as400.SSLConfiguration object.

Note: The previously mentioned restrictions mentioned may change in future releases, so that the native iSeries JSSE provider could allow you to plug in non-native components (for example, JKS KeyStore or IbmX509 TrustManagerFactory).



Examples: IBM Java Secure Sockets Extension

The JSSE examples show how a client and a server can use the native iSeries JSSE provider to create a context that enables secure communications.

Note: Both examples use the native iSeries JSSE provider, regardless of the properties specified by the java.security file.

Example: SSL Client Using an SSLContext Object

This example client program utilizes an SSLContext object, which it initializes to use the "MY_CLIENT_APP" application ID. This program will use the native iSeries implementation regardless of what is specified in the java.security file.

Example: SSL Server Using an SSLContext Object

The following server program utilizes an SSLContext object that it initializes with a previously created keystore file. The keystore file has a name of /home/keystore.file and a keystore password of password.

The example program needs the keystore file in order to create an IbmIseriesKeyStore object. The KeyStore object must specify MY_SERVER_APP as the application identifier.

To create the keystore file, you can use either of the following commands:

- From a Qshell command prompt:

```
java com.ibm.as400.SSLConfiguration -create -keystore /home/keystore.file
-storepass password -appid MY_SERVER_APP
```

For more information about using Java commands with Qshell, see the following topic:

Qshell

- From an iSeries command prompt:

```
RUNJAVA CLASS(com.ibm.as400.SSLConfiguration) PARM('-create' '-keystore'
'/home/keystore.file' '-storepass' 'password' '-appid' 'MY_SERVER_APP')
```

The following disclaimer applies to all of the IBM JSSE examples:

Code example disclaimer

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

All sample code is provided by IBM for illustrative purposes only. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

All programs contained herein are provided to you "AS IS" without any warranties of any kind. The implied warranties of non-infringement, merchantability and fitness for a particular purpose are expressly disclaimed.



Example: SSL client using an SSLContext object

Note: Read the Code example disclaimer for important legal information.

```
////////////////////////////////////
//
// This example client program utilizes an SSLContext object, which it initializes
// to use the "MY_CLIENT_APP" application ID.
//
// The example uses the native iSeries JSSE provider, regardless of the
// properties specified by the java.security file.
//
// Command syntax:
//   java -Djava.version=1.4 SslClient
//
// Note that "-Djava.version=1.4" is unnecessary when you have configured
// J2SDK version 1.to be used by default.
//
////////////////////////////////////

import java.io.*;
import javax.net.ssl.*;

/**
 * SSL Client Program.
 */
public class SslClient {

    /**
     * SslClient main method.
     *
     * @param args the command line arguments (not used)
     */
    public static void main(String args[]) {
        /**
         * Set up to catch any exceptions thrown.
         */
        try {
            /**
             * Initialize an SSLConfiguration object to specify an application
             * ID. "MY_CLIENT_APP" must be registered and configured
             * correctly with the Digital Certificate Manager (DCM).
             */
            SSLConfiguration config = new SSLConfiguration();
            config.setApplicationId("MY_CLIENT_APP");
            /**
             * Get a KeyStore object from the SSLConfiguration object.
             */
            Char[] password = "password".toCharArray();
            KeyStore ks = config.getKeyStore(password);
            /**
             * Allocate and initialize a KeyManagerFactory.
             */
            KeyManagerFactory kmf =
                KeyManagerFactory.getInstance("IbmISeriesX509");
            KmF.init(ks, password);
            /**
             * Allocate and initialize a TrustManagerFactory.
             */
            TrustManagerFactory tmf =
```



```

        TrustManagerFactory.getInstance("IbmISeriesX509");
tmf.init(ks);
/*
 * Allocate and initialize an SSLContext.
 */
SSLContext c =
    SSLContext.getInstance("SSL", "quot;);
C.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);
/*
 * Get the an SSLSocketFactory from the SSLContext.
 */
SSLSocketFactory sf = c.getSocketFactory();
/*
 * Create an SSLSocket.
 *
 * Change the hard-coded IP address to the IP address or host name
 * of the server.
 */
SSLSocket s = (SSLSocket) sf.createSocket("1.1.1.1", 13333);
/*
 * Send a message to the server using the secure session.
 */
String sent = "Test of java SSL write";
OutputStream os = s.getOutputStream();
os.write(sent.getBytes());
/*
 * Write results to screen.
 */
System.out.println("Wrote " + sent.length() + " bytes...");
System.out.println(sent);
/*
 * Receive a message from the server using the secure session.
 */
InputStream is = s.getInputStream();
byte[] buffer = new byte[1024];
int bytesRead = is.read(buffer);
if (bytesRead == -1)
    throw new IOException("Unexpected End-of-file Received");
String received = new String(buffer, 0, bytesRead);
/*
 * Write results to screen.
 */
System.out.println("Read " + received.length() + " bytes...");
System.out.println(received);
} catch (Exception e) {
    System.out.println("Unexpected exception caught: " +
        e.getMessage());
    e.printStackTrace();
}
}
}

```



Example: SSL server using an SSLContext object

Note: Read the Code example disclaimer for important legal information.

```

////////////////////////////////////
//
// The following server program utilizes an SSLContext object that it
// initializes with a previously created keystore file.
//
// The keystore file has the following name and keystore password:
//   File name: /home/keystore.file
//   Password: password

```

```

//
// The example program needs the keystore file in order to create an
// IbmISeriesKeyStore object. The KeyStore object must specify MY_SERVER_APP as
// the application identifier.
//
// To create the keystore file, you can use the following Qshell command:
//
// java com.ibm.as400.SSLConfiguration -create -keystore /home/keystore.file
// -storepass password -appid MY_SERVER_APP
//
// Command syntax:
// java -Djava.version=1.4 JavaSslServer
//
// Note that "-Djava.version=1.4" is unnecessary when you have configured
// J2SDK version 1.to be used by default.
//
//
//

```

```

import java.io.*;
import javax.net.ssl.*;

```

```

/**
 * Java SSL Server Program using Application ID.
 */
public class JavaSslServer {

    /**
     * JavaSslServer main method.
     *
     * @param args the command line arguments (not used)
     */
    public static void main(String args[]) {
        /*
         * Set up to catch any exceptions thrown.
         */
        try {
            /*
             * Allocate and initialize a KeyStore object.
             */
            Char[] password = "password".toCharArray();
            KeyStore ks = KeyStore.getInstance("IbmISeriesKeyStore");
            FileInputStream fis = new FileInputStream("/home/keystore.file");
            Ks.load(fis, password);
            /*
             * Allocate and initialize a KeyManagerFactory.
             */
            KeyManagerFactory kmf =
                KeyManagerFactory.getInstance("IbmISeriesX509");
            KmF.init(ks, password);
            /*
             * Allocate and initialize a TrustManagerFactory.
             */
            TrustManagerFactory tmf =
                TrustManagerFactory.getInstance("IbmISeriesX509");
            tmf.init(ks);
            /*
             * Allocate and initialize an SSLContext.
             */
            SSLContext c =
                SSLContext.getInstance("SSL", "IbmISeriesSslProvider");
            C.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);
            /*
             * Get the an SSLServerSocketFactory from the SSLContext.
             */
            SSLServerSocketFactory sf = c.getSSLServerSocketFactory();
            /*
             * Create an SSLServerSocket.

```

```

    */
    SSLServerSocket ss =
        (SSLServerSocket) sf.createServerSocket(13333);
    /*
    * Perform an accept() to create an SSLSocket.
    */
    SSLSocket s = (SSLSocket) ss.accept();
    /*
    * Receive a message from the client using the secure session.
    */
    InputStream is = s.getInputStream();
    byte[] buffer = new byte[1024];
    int bytesRead = is.read(buffer);
    if (bytesRead == -1)
        throw new IOException("Unexpected End-of-file Received");
    String received = new String(buffer, 0, bytesRead);
    /*
    * Write results to screen.
    */
    System.out.println("Read " + received.length() + " bytes...");
    System.out.println(received);
    /*
    * Echo the message back to the client using the secure session.
    */
    OutputStream os = s.getOutputStream();
    os.write(received.getBytes());
    /*
    * Write results to screen.
    */
    System.out.println("Wrote " + received.length() + " bytes...");
    System.out.println(received);
} catch (Exception e) {
    System.out.println("Unexpected exception caught: " +
        e.getMessage());
    e.printStackTrace();
}
}
}

```

Chapter 3. Java Authentication and Authorization Service

The Java^(TM) Authentication and Authorization Service (JAAS) is a standard extension to the Java 2 Software Development Kit (J2SDK), Standard Edition. Currently, J2SDK provides access controls that are based on where the code originated and who signed the code (code source-based access controls). It lacks, however, the ability to enforce additional access controls based on who runs the code. JAAS provides a framework that adds this support to the Java 2 security model.

The JAAS API is used by IBM and Sun Microsystems, Inc. as an extension to the J2SDK, [»](#) versions 1.2 and 1.3. [«](#) IBM and Sun are introducing this extension to allow the association of a specific user or identity to the current Java thread. This is done by using `javax.security.auth.Subject` methods and, optionally, with the underlying operating system thread using `com.ibm.security.auth.ThreadSubject` methods.

[»](#) **Note:** For J2SDK, version 1.4, JAAS is no longer an extension, but is part of the base SDK. [«](#)

The JAAS implementation on the iSeries server is compatible with the implementation of Sun Microsystems, Inc. This documentation covers the unique aspects of the iSeries implementation. We assume that you are familiar with the general documentation for the JAAS extensions. To make it easier for you to work with that and our iSeries information, we provide the following links.

- The API Developers Guide provides information on using the JAAS API in software development.
- Login/Authentication Module Developers Guide focuses on the authentication aspects of JAAS.
- JAAS API Specification contains the Javadoc information on JAAS.

Select any of these topics for more details about how to use JAAS:

- Prepare and configure the iSeries server for JAAS
- JAAS samples
- iSeries-server specific JAAS Javadoc

Prepare and configure an iSeries server for Java Authentication and Authorization Service

You must meet software requirements and configure your iSeries server to use Java^(TM) Authentication and Authorization Service (JAAS).

Software requirements to run JAAS 1.0 on an iSeries server

Install the following licensed programs:

- Java 2 SDK, version 1.4 (J2SDK)
- The IBM Toolbox for Java (mod 4) Licensed Program (5722-JC1) is required to change the OS thread identity. It contains the `ProfileTokenCredential` classes needed to support the changing of iSeries OS thread identity and the native implementation classes.

Configure the system

To configure the system to use JAAS, follow these steps:

1. For JDKs 1.2 and 1.3, add a symbolic link to the extension directory for the `jaas13.jar` file. The extension class loader should load the JAR file. Run this command (all one line) on the iSeries command line to add the link:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/Java400/ext/jaas13.jar')
NEWLNK('/QIBM/ProdData/Java400/jdk13/lib/ext/jaas13.jar')
```

» **Note:** For JDK 1.4, you do not need to add a symbolic link to the extension directory. JAAS is part of the base SDK for this version. «

2. A default login.config file is provided in `{java.home}/lib/security` which invokes `com.ibm.as400.security.auth.login.BasicAuthenticationLoginModule`. This login.config file attaches a single use `ProfileTokenCredential` to the authenticated subject. If you want to use your own login.config file with different options, you may include the following system property when invoking your application:

```
-Djava.security.auth.login.config=your login.config file
```

3. Add a symbolic link to the extension directory for the `jt400Native.jar` file. This allows the extension class loader to load this file. The `jaas13.jar` file requires this JAR file for the credential implementation classes that are part of the IBM Toolbox for Java. The application class loader can also load this file by including it in the CLASSPATH. If this file is loaded from the class path directory, do not add the symbolic link to the extension directory.

» Symbolically linking the `jt400Native.jar` file to the `/QIBM/ProdData/Java400/jdk14/lib/ext` directory forces all JDK 1.4 users on the server to run with this version of `jt400Native.jar`. « This may not be desirable if various users require different versions of the IBM Toolbox for Java classes. Other options include putting `jt400Native.jar` in the application CLASSPATH as described previously. Another option is to add the symbolic link to your own directory and then include that directory in the extension directory classpath by specifying the `java.ext.dirs` system property when invoking the application.

To link the `jt400Native.jar` file to the `/QIBM/ProdData/Java400/jdk13/lib/ext` directory, run this command on the iSeries command line to add the link:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/jt400/lib/jt400Native.jar')
NEWLNK('/QIBM/ProdData/Java400/jdk13/lib/ext/jt400Native.jar')
```

» To link the `jt400Native.jar` file to the `/QIBM/ProdData/Java400/jdk14/lib/ext` directory, run this command on the iSeries command line to add the link:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/jt400/lib/jt400Native.jar')
NEWLNK('/QIBM/ProdData/Java400/jdk14/lib/ext/jt400Native.jar')
```

«

To link the `jt400Native.jar` file to your own directory, do the following:

- a. Run this command on the iSeries command line to add the link:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/jt400/lib/jt400Native.jar')
NEWLNK('your extension directory/jt400Native.jar')
```

- b. When calling your java program, use the following pattern:

```
java -Djava.ext.dirs=your extension directory:default
extension directories
```

Note: See the IBM Toolbox for Java for information on the iSeries credential classes. Click on **Security classes**. Click on **Authentication Services**. Click on **ProfileTokenCredential** class. Click on **Package**.

4. Update the Java 2 policy files to grant the appropriate permissions to the actual locations of the IBM Toolbox for Java JAR files. Even though these files may be symbolically linked to the extension directories and those directories are granted `java.security.AllPermission` in the `{java.home}/lib/security/java.policy` file, authorization is based on the actual location of the JAR files. To successfully use the credential classes in the IBM Toolbox for Java, add the following to the Java 2 policy file of your application:

```
grant codeBase "file:/QIBM/ProdData/OS400/jt400/lib/jt400Native.jar"
{
    permission javax.security.auth.AuthPermission "modifyThreadIdentity";
}
```

```
permission java.lang.RuntimePermission "loadLibrary.*";
permission java.lang.RuntimePermission "writeFileDescriptor";
permission java.lang.RuntimePermission "readFileDescriptor";
}
```

You also need to add these permissions for the codeBase of your application since the operations performed by the IBM Toolbox for Java JAR files do not run in privileged mode.

See the API Developers Guide for information on the Java 2 policy files.

5. Make sure the iSeries Host Servers are started and running. The ProfileTokenCredential classes that reside in the Toolbox, for example, jt400Native.jar, are used as the credentials that are attached to the authenticated subject. The credential classes require access to the Host Servers. You can verify that the servers are started and running by typing the following on the iSeries command prompt:

- StrHostSVR *all
- StrTcpSvr *DDM

If the servers have already been started, these steps do nothing. If the servers are not started, they are started.

Java Authentication and Authorization Service samples

In this information, we provide a link to some samples of Java^(TM) Authentication and Authorization Service (JAAS) on an iSeries server. There are two JAAS samples included with the documentation, HelloWorld and SampleThreadSubjectLogin. Click on these links for instructions and source code:

- [HelloWorld](#)
- [SampleThreadSubjectLogin](#)



Chapter 4. IBM Java Generic Security Service (JGSS)

The Java Generic Security Service (JGSS) provides a generic interface for authentication and secure messaging. Under this interface you can plug a variety of security mechanisms based on secret-key, public-key, or other security technologies.

By abstracting the complexity and peculiarities of the underlying security mechanisms to a standardized interface, JGSS provides the following benefits to the development of secure networking applications:

- You can develop the application to a single abstract interface
- You can use the application with different security mechanisms without any changes

JGSS defines the Java bindings for the Generic Security Service Application Programming Interface (GSS-API), which is a cryptographic API that has been standardized by the Internet Engineering Task Force (IETF) and adopted by the X/Open Group.

The IBM implementation of JGSS is called IBM JGSS. IBM JGSS is an implementation of the GSS-API framework that uses Kerberos V5 as the default underlying security system. It also features a Java^(TM) Authentication and Authorization Service (JAAS) login module for creating and using Kerberos credentials. In addition, you can have JGSS perform JAAS authorization checks when you use those credentials.

IBM JGSS includes a native iSeries JGSS provider, a Java JGSS provider, and Java versions of the Kerberos credential management tools (kinit, ktab, and klist).

Note: The native iSeries JGSS provider uses the native iSeries Network Authentication Services (NAS) library. When you use the native provider, you must use the native iSeries Kerberos utilities. For more information, see JGSS providers .

For more information about using JGSS, see the following topics:

JGSS concepts

Introduces JGSS concepts, including a high-level description of GSS-API operations and a brief discussion of security mechanisms.

Configuring your server to use JGSS

Find out how to configure your iSeries server to use IBM JGSS with the Java^(TM) 2 Software Development Kit, Standard Edition (J2SDK). Information includes identifying and setting the permissions necessary to use JGSS with a security manager.

Running JGSS applications

Learn about running JGSS applications on your iSeries servers. Documentation includes an explanation of operational concepts and instructions for using JAAS.

Developing JGSS applications

See how to use JGSS to develop secure applications. Learn about generating transport tokens, creating JGSS objects, establishing context, and more.





JGSS javadoc reference information

Review the javadoc information for classes and methods in the org.ietf.jgss api package and for the Java versions of Kerberos credential management tools (kinit, ktab, and klist).

JGSS samples

Use the sample programs to discover how you can use JGSS in your applications. The sample documentation includes Java source code, instructions for running the samples, configuration and policy files, and more.

To read more about Java security and the generic security service, see the following documentation:

- J2SDK Security enhancement  from Sun Microsystems, Inc. that contains links to more Java GSS-API information
- Internet Engineering Task Force (IETF) RFC 2743 Generic Security Services Application Programming Interface Version 2, Update 1 
- IETF RFC 2853 Generic Security Service API Version 2: Java Bindings 
- The X/Open Group GSS-API Extensions for DCE 

Note: Read the Code example disclaimer for important legal information.  


JGSS concepts

JGSS operations consist of four distinct stages, as standardized by the Generic Security Service Application Programming Interface (GSS-API):

1. Gathering of credentials for principals
2. Creating and establishing a security context between the communicating peer principals
3. Exchanging secure messages between the peers
4. Cleaning up and releasing resources

Additionally, JGSS leverages the Java Cryptographic Architecture to offer seamless pluggability of different security mechanisms.

Use the following links to read high-level descriptions of these important JGSS concepts.

- Principals and credentials
- Context establishment
- Message protection and exchange
- Resource cleanup and release
- Security mechanisms 




Principals and credentials

The identity under which an application engages in JGSS secure communication with a peer is called a principal. A principal may be a real user or an unattended service. A principal acquires security mechanism-specific credentials as proof of identity under that mechanism. For example, when using the Kerberos mechanism, a principal's credential is in the form of a ticket-granting ticket (TGT) issued by a Kerberos key distribution center (KDC). In a multi-mechanism environment, a GSS-API credential can contain multiple credential elements, each element representing an underlying mechanism credential.

The GSS-API standard does not prescribe how a principal acquires credentials, and GSS-API implementations typically do not provide a means for credential acquisition. A principal obtains credentials before using GSS-API; GSS-API merely queries the security mechanism for credentials on behalf of the principal.

IBM JGSS includes Java versions of Kerberos credential management tools kinit, ktab and klist. Additionally, IBM JGSS enhances the standard GSS-API by providing an optional Kerberos login interface that uses JAAS. The pure Java JGSS provider supports the optional login interface; the native iSeries provider does not. For more information, see the following topics:

- Obtaining Kerberos credentials
- JGSS providers 



Context establishment

Having acquired security credentials, the two communicating peers establish a security context using their credentials. Although the peers establish a single joint context, each peer maintains its own local copy of the context. Context establishment involves the initiating peer authenticating itself to the accepting peer. The initiator optionally may request mutual authentication, in which case the acceptor authenticates itself to the initiator.

When context establishment is complete, the established context embodies state information (such as shared cryptographic keys) that enable subsequent secure message exchange between the two peers. <<



Message protection and exchange

Following context establishment, the two peers are ready to engage in secure message exchanges. The originator of the message calls on its local GSS-API implementation to encode the message, which ensures message integrity and, optionally, message confidentiality. The application then transports the resulting token to the peer.

The local GSS-API implementation of the peer uses information from the established context in the following ways:

- Verifies the integrity of the message
- Deciphers the message, if the message was encrypted <<



Resource cleanup and release

In order to free up resources, a JGSS application deletes a context that is no longer needed. Although a JGSS application can access a deleted context, any attempt to use it for message exchange results in an exception. << >>

Security mechanisms



The GSS-API consists of an abstract framework over one or more underlying security mechanisms. How the framework interacts with the underlying security mechanisms is implementation specific. Such implementations exist in two general categories:

- At one extreme, a monolithic implementation tightly binds the framework to a single mechanism. This kind of implementation precludes the use of other mechanisms or even different implementations of the same mechanism.
- At the other end of the spectrum, a highly modular implementation offers ease of use and flexibility. This kind of implementation offers the ability to seamlessly and easily plug different security mechanisms and their implementations into the framework.

IBM JGSS falls into the latter category. As a modular implementation, IBM JGSS leverages the provider framework defined by the Java Cryptographic Architecture (JCA) and treats any underlying mechanism as a (JCA) provider. A JGSS provider supplies a concrete implementation of a JGSS security mechanism. An application can instantiate and use multiple mechanisms.


It is possible for a provider to support multiple mechanisms, and JGSS makes it easy to use different security mechanisms. However, the GSS-API does not provide a means for two communicating peers to choose a mechanism when multiple mechanisms are available. One way to choose a mechanism is to

start with the Simple And Protected GSS-API Negotiating Mechanism (SPNEGO), a pseudo-mechanism that negotiates an actual mechanism between the two peers. IBM JGSS does not include a SPNEGO mechanism.

For more information about SPNEGO, see Internet Engineering Task Force (IETF) RFC 2478 The Simple and Protected GSS-API Negotiation Mechanism  .  

Configuring your iSeries server to use IBM JGSS

How you configure your iSeries server to use JGSS depends on which version of the Java 2 Software Development Kit (J2SDK) that you run on your server. For more information about configuring your iSeries server to use JGSS, use the following links:

- Using JGSS with J2SDK, version 1.3
- Using JGSS with J2SDK, version 1.4
- Configuring JGSS to use the native iSeries JGSS provider 



Configuring your iSeries server to use JGSS with J2SDK, version 1.3

When you use the Java 2 Software Development Kit (J2SDK), version 1.3 on your iSeries server, you need to prepare and configure your server to use JGSS. The default configuration uses the pure Java JGSS provider.

Software requirements

To use JGSS with J2SDK, version 1.3, your server must have Java Authentication and Authorization Service (JAAS) 1.3 installed.

Configuring your server to use JGSS

To configure your server to use JGSS with J2SDK, version 1.3, add a symbolic link to the extension directory for the `ibmjgssprovider.jar` file. The `ibmjgssprovider.jar` file contains the JGSS classes and the pure Java JGSS provider. Adding the symbolic link enables the extension class loader to load the `ibmjgssprovider.jar` file.

Adding the symbolic link

To add the symbolic link, on an iSeries command line, type the following command (on a single line) and press **ENTER**:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/Java400/ext/ibmjgssprovider.jar')
NEVLNK('/QIBM/ProdData/Java400/jdk13/lib/ext/ibmjgssprovider.jar')
```



Note: The default Java 1.3 policy on the iSeries server grants the appropriate permissions to JGSS. If you plan to create your own `java.policy` file, see JVM permissions for a list of permissions to grant `ibmjgssprovider.jar`.

Changing JGSS providers

After configuring your server to use JGSS, which uses the pure Java provider as the default, you can configure JGSS to use the native iSeries JGSS provider. Then, after you configure JGSS to use the native provider, you can easily switch between the two providers. For more information, see the following topics:

- JGSS providers
- Configuring JGSS to use the native iSeries JGSS provider

Security managers

If you are running your IBM JGSS application with a Java security manager enabled, see Using a security manager.  

Configuring JGSS to use the native iSeries JGSS provider

IBM JGSS uses the pure Java provider by default. You have the option to use the native iSeries JGSS provider. For more information about the different providers, see JGSS providers.

Software requirements

The native iSeries JGSS provider must be able to access classes in IBM Toolbox for Java. For instructions about how to access Toolbox for Java, see Enabling the native iSeries JGSS provider to access IBM Toolbox for Java.

Make sure that you have configured the network authentication service. For more information, see Network authentication service.

Specifying the native iSeries JGSS provider

Before you can use the native iSeries JGSS provider with J2SDK, version 1.3, ensure that you have configured your server to use JGSS. For more information, see Configuring your iSeries server to use JGSS with J2SDK, version 1.3. If you are using J2SDK, version 1.4, JGSS is already configured.

Note: In the following instructions, `{java.home}` denotes the path to the location of the version of Java that you are using on your server. For example, if you are using J2SDK, version 1.4, `{java.home}` is `/QIBM/ProdData/Java400/jdk14`. Remember to replace `{java.home}` in the commands with the actual path to the Java home directory.

To configure JGSS to use the native iSeries JGSS provider, complete the following tasks:

- Add a symbolic link to the extension directory for the native iSeries provider JAR file
- Add the native iSeries JGSS provider to the security provider list in the `java.security` file

Adding a symbolic link

To add a symbolic link to the extension directory for the `ibmjgssiseriesprovider.jar` file, on an iSeries command line, type the following command (on a single line) and press **ENTER**:

```
ADDLNK OBJ('/QIBM/ProdData/OS400/Java400/ext/ibmjgssiseriesprovider.jar')
NEWLNK('{java.home}/lib/ext/ibmjgssiseriesprovider.jar')
```

After you add a symbolic link to the extension directory for the `ibmjgssiseriesprovider.jar` file, the extension class loader will load the JAR file.

Adding the provider to the security provider list

Add the native provider to the the security provider list in the `java.security` file.

1. Open `{java.home}/lib/security/java.security` for editing.
2. Find the security provider list. It should be near the top of the `java.security` file and should look something like:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.sun.rsajca.Provider
security.provider.3=com.ibm.crypto.provider.IBMJCE
security.provider.4=com.ibm.security.jgss.IBMJGSSProvider
```

3. Add the native iSeries JGSS provider to the security provider list before the original Java provider. In other words, add `com.ibm.iseries.security.jgss.IBMJGSSiSeriesProvider` to the list with a lower number than `com.ibm.jgss.IBMJGSSProvider`, then update the position of `IBMJGSSProvider`. For example:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.sun.rsajca.Provider
security.provider.3=com.ibm.crypto.provider.IBMJCE
security.provider.4=com.ibm.iseries.security.jgss.IBMJGSSiSeriesProvider
security.provider.5=com.ibm.security.jgss.IBMJGSSProvider
```

Notice that the IBMJGSSiSeriesProvider became the fourth entry in the list and IBMJGSSProvider became the fifth entry. Also, check that entry numbers in the security provider list are sequential and that each entry increments the entry number by only one.

4. Save and close the java.security file. <<



Configuring your iSeries server to use JGSS with J2SDK, version 1.4

When you use the Java 2 Software Development Kit (J2SDK), version 1.4 on your iSeries server, JGSS is already configured. The default configuration uses the pure Java JGSS provider.

Changing JGSS providers

You can configure JGSS to use the native iSeries JGSS provider instead of the pure Java JGSS provider. Then, after you configure JGSS to use the native provider, you can easily switch between the two providers. For more information, see the following topics:

- JGSS providers
- Configuring JGSS to use the native iSeries JGSS provider

Security managers

If you are running your JGSS application with a Java security manager enabled, see Using a security manager. << >>

JGSS providers

IBM JGSS includes a native iSeries JGSS provider and a pure Java JGSS provider. The provider that you choose to use depends on the needs of your application.

The pure Java JGSS provider offers the following features:

- Ensures the greatest level of portability for your application
- Works with the optional JAAS Kerberos login interface
- Compatible with the Java Kerberos credential management tools

The native iSeries JGSS provider offers the following features:

- Uses the native iSeries Kerberos libraries
- Compatible with Qshell Kerberos credential management tools
- JGSS applications run faster

Note: Both JGSS providers adhere to the GSS-API specification and so are compatible with each other. In other words, an application that uses the pure Java JGSS provider can interoperate with an application that uses the native iSeries JGSS provider.

Changing the JGSS provider

Note: If your server is running J2SDK, version 1.3, before changing to the native iSeries JGSS provider, make sure that you have configured your server to use JGSS. For more information, see the following topics:

- Configuring your iSeries server to use JGSS with J2SDK, version 1.3
- Configuring JGSS to use the Native JGSS Provider

You can easily change the JGSS provider by using one of the following options:

- Edit the security provider list in `${java.home}/lib/security/java.security`

Note: `${java.home}` denotes the path to the location of the version of Java that you are using on your server. For example, if you are using J2SDK, version 1.3, `${java.home}` is `/QIBM/ProdData/Java400/jdk13`.

- Specify the provider name in your JGSS application by using either `GSSManager.addProviderAtFront()` or `GSSManager.addProviderAtEnd()`. For more information, see the `GSSManager` javadoc. [↩](#)



Using a security manager

If you are running your JGSS application with a Java security manager enabled, you need to ensure that your application and JGSS have the necessary permissions. For more information about the permissions you need to use JGSS, see the following topics:

- JVM permissions
- JAAS permission checks [↩](#)



JVM permissions

In addition to the access control checks performed by JGSS, the Java virtual machine (JVM) performs authorization checks when accessing a variety of resources, including files, Java properties, packages, and sockets.

For more information about using JVM permissions, see [Permissions in the Java 2 SDK](#) .

The following list identifies the permissions required when you use the JAAS features of JGSS or use JGSS with a security manager:

- `javax.security.auth.AuthPermission "modifyPrincipals"`
- `javax.security.auth.AuthPermission "modifyPrivateCredentials"`
- `javax.security.auth.AuthPermission "getSubject"`
- `javax.security.auth.PrivateCredentialPermission "javax.security.auth.kerberos.KerberosKey javax.security.auth.kerberos.KerberosPrincipal **", "read"`
- `javax.security.auth.PrivateCredentialPermission "javax.security.auth.kerberos.KerberosTicket javax.security.auth.kerberos.KerberosPrincipal **", "read"`
- `java.util.PropertyPermission "com.ibm.security.jgss.debug", "read"`
- `java.util.PropertyPermission "DEBUG", "read"`
- `java.util.PropertyPermission "java.home", "read"`
- `java.util.PropertyPermission "java.security.krb5.conf", "read"`
- `java.util.PropertyPermission "java.security.krb5.kdc", "read"`
- `java.util.PropertyPermission "java.security.krb5.realm", "read"`
- `java.util.PropertyPermission "javax.security.auth.useSubjectCredsOnly", "read"`
- `java.util.PropertyPermission "user.dir", "read"`
- `java.util.PropertyPermission "user.home", "read"`
- `java.lang.RuntimePermission "accessClassInPackage.sun.security.action"`
- `java.security.SecurityPermission "putProviderProperty.IBMJGSSProvider" ↩`



JAAS permission checks

IBM JGSS performs runtime permission checks at the time the JAAS-enabled program uses credentials and accesses services. You can disable this optional JAAS feature by setting the Java property `javax.security.auth.useSubjectCredsOnly` to `false`. Moreover, JGSS performs permission checks only when the application runs with a security manager.

JGSS performs permission checks against the Java policy that is in effect for the current access control context. JGSS performs the following specific permission checks:

- `javax.security.auth.kerberos.DelegationPermission`
- `javax.security.auth.kerberos.ServicePermission`

DelegationPermission check

The `DelegationPermission` allows the security policy to control the use of the ticket forwarding and proxying features of Kerberos. Using these features, a client can allow a service to act on behalf of the client.

`DelegationPermission` takes two arguments, in the following order:

1. The subordinate principal, which is the name of the service principal that acts on behalf of, and under the authority of, the client
2. The name of the service that the client wants to allow the subordinate principal to use

Example: Using the DelegationPermission check


In the following example, `superSecureServer` is the subordinate principal and `krbtgt/REALM.IBM.COM@REALM.IBM.COM` is the service that we want to allow `superSecureServer` to use on behalf of the client. In this case, the service is the ticket-granting ticket for the client, which means that `superSecureServer` can get a ticket for any service on behalf of the client.

```
permission javax.security.auth.kerberos.DelegationPermission
    "\"superSecureServer/host.ibm.com@REALM.IBM.COM\"
    \"krbtgt/REALM.IBM.COM@REALM.IBM.COM\"";
```

In the previous example, `DelegationPermission` grants the client permission to get a new ticket-granting ticket from the Key Distribution Center (KDC) that only `superSecureServer` can use. After the client has sent the new ticket-granting ticket to `superSecureServer`, `superSecureServer` has the ability to act on behalf of the client.

The following example enables the client to get a new ticket that allows `superSecureServer` to access only the ftp service on behalf of the client:

```
permission javax.security.auth.kerberos.DelegationPermission
    "\"superSecureServer/host.ibm.com@REALM.IBM.COM\"
    \"ftp/ftp.ibm.com@REALM.IBM.COM\"";
```

For more information, see the `javax.security.auth.kerberos.DelegationPermission` class in the J2SDK documentation  on the Sun Web site.

ServicePermission check

`ServicePermission` checks restrict the use of credentials for context initiation and acceptance. A context initiator must have permission to initiate a context. Likewise, a context acceptor must have permission to accept a context.




Example: Using the ServicePermission check

The following example allows the client side to initiate a context with the ftp service by granting permission to the client:


```
permission javax.security.auth.kerberos.ServicePermission
    "ftp/host.ibm.com@REALM.IBM.COM", "initiate";
```

The following example allows the server side to access and use the secret key for the ftp service by granting permission to the server:

```
permission javax.security.auth.kerberos.ServicePermission
    "ftp/host.ibm.com@REALM.IBM.COM", "accept";
```

For more information, see the `javax.security.auth.kerberos.ServicePermission` class in the J2SDK documentation  on the Sun Web site.  

Running IBM JGSS applications

The IBM Java Generic Security Service (JGSS) API 1.0 shields secure applications from the complexities and peculiarities of the different underlying security mechanisms. JGSS uses features provided by Java Authentication and Authorization Service (JAAS) and IBM Java Cryptography Extension (JCE).

JGSS features include:

- Identity authentication
- Message integrity and confidentiality
- Optional JAAS Kerberos login interface and authorization checks

For more information about running JGSS applications, see the following topics:

Obtaining Kerberos credentials

Find out how to obtain Kerberos credentials and create secret keys. Learn about using JAAS to perform Kerberos logins and authorization checks and review a list of JAAS permissions required by the Java virtual machine (JVM).


Configuration and policy files

Learn about the different kinds of supporting files you need to run JGSS, including configuration files, policy files, the Java master security properties file, and the credentials cache.

Debugging

Read about using JGSS debugging to categorize and display helpful debugging messages.


JGSS samples

Use the sample programs to test and verify your JGSS setup. The sample documentation includes Java source code, instructions for running the samples, configuration and policy files, and more. 



Obtaining Kerberos credentials and creating secret keys

The GSS-API does not define a way to get credentials. For this reason, the IBM JGSS Kerberos mechanism requires that the user obtain Kerberos credentials by using one of the following methods:

- Kinit and Ktab tools
- Optional JAAS Kerberos login interface 



The Kinit and Ktab tools

Your choice of a JGSS provider determines which tools that you use to obtain Kerberos credentials and secret keys.

Using the pure Java JGSS provider

If you are using the pure Java JGSS provider, use the IBM JGSS Kinit and Ktab tools to obtain credentials and secret keys. The Kinit and Ktab tools use command-line interfaces and provide options similar to those offered by other versions.

- You can obtain Kerberos credentials by using the Kinit tool. This tool contacts the Kerberos Distribution Center (KDC) and obtains a ticket-granting ticket (TGT). The TGT allows you to access other Kerberos-enabled services, including those that use the GSS-API.
- A server can obtain a secret key by using the Ktab tool. JGSS stores the secret key in the key table file on the server. See the Ktab Java documentation for more information.

Alternatively, your application can use the JAAS Login interface to obtain TGTs and secret keys. For more information, see the following:

- Kinit javadoc
- Ktab javadoc
- JAAS login interface.

Using the native iSeries JGSS provider

If you are using the native iSeries JGSS provider, use the Qshell kinit and klist utilities. For more information, see Utilities for Kerberos credentials and key tables. [«](#) [»](#)

JAAS Kerberos login interface

IBM JGSS features a Java Authentication and Authorization Service (JAAS) Kerberos login interface. You can disable this feature by setting the Java property `javax.security.auth.useSubjectCredsOnly` to `false`.

Note: Although the pure Java JGSS provider can use the login interface, the native iSeries JGSS provider can not.

For more information about JAAS, see Java Authentication and Authorization Service.

JAAS and JVM permissions

If you are using a security manager, you need to ensure that your application and JGSS have the necessary JVM and JAAS permissions. For more information, see Using a security manager.

JAAS configuration file options

The login interface requires a JAAS configuration file that specifies `com.ibm.security.auth.module.Krb5LoginModule` as the login module to be used. The following table lists the options that `Krb5LoginModule` supports. Note that the options are not case-sensitive.

Option Name	Value	Default	Explanation
<code>principal</code>	<code><string></code>	None; prompted for.	Kerberos principal name
<code>credsType</code>	<code>initiator acceptor both</code>	<code>initiator</code>	The JGSS credential type
<code>forwardable</code>	<code>true false</code>	<code>false</code>	Whether to acquire a forwardable ticket-granting ticket (TGT)
<code>proxiable</code>	<code>true false</code>	<code>false</code>	Whether to acquire a proxiable TGT
<code>useCcache</code>	<code><URL></code>	Don't use ccache	Retrieve TGT from the specified credential cache
<code>useKeytab</code>	<code><URL></code>	Don't use key table	Retrieve secret key from the specified key table
<code>useDefaultCcache</code>	<code>true false</code>	Don't use default ccache	Retrieve TGT from default credential cache
<code>useDefaultKeytab</code>	<code>true false</code>	Don't use default key table	Retrieve secret key from the specified key table

For a simple example of using Krb5LoginModule, see the sample JAAS login configuration file.

Option incompatibilities

Some Krb5LoginModule options, excluding principal name, are incompatible with each other, meaning that you can not specify them together. The following table represents compatible and incompatible login module options.

Indicators in the table describe the relationship between the two associated options:

- X = Incompatible
- N/A = Inapplicable combination
- Blank = Compatible

Krb5LoginModule option	credsType=initiator	credsType=acceptor	credsType=both	forward	proxy	useCcache	useKeytab	useDefaultCcache	useDefaultKeytab
credsType=initiator		N/A	N/A				X		X
credsType=acceptor	N/A		N/A	X	X	X		X	
credsType=both	N/A	N/A							
forwardable		X				X	X	X	X
proxiabile		X				X	X	X	X
useCcache		X		X	X		X	X	X
useKeytab	X			X	X	X		X	X
useDefaultCcache		X		X	X	X	X		X
useDefaultKeytab	X			X	X	X	X	X	

Principal name option

You can specify a principal name in combination with any other option. When you do not specify a principal name, the Krb5LoginModule may prompt the user for a principal name. Whether or not Krb5LoginModule prompts the user depends on the other options that you specify. For more information, see Prompting for the principal name and password.

Service principal name format

You must use one of the following formats to specify a service principal name:

- <service_name> (for example, superSecureServer)
- <service_name>@<host> (for example, superSecureServer@myhost)

In the latter format, <host> is the hostname of the machine on which the service resides. You can (but do not have to) use a fully qualified hostname.

Note: JAAS recognizes certain characters as delimiters. When you use any of the following characters in a JAAS string (such as a principal name), enclose the character in quotes:

- _ (underscore)
- :
- / (forward slash)
- \ (back slash)

Prompting for the principal name and password

The options that you specify in the JAAS configuration file determine whether the Krb5LoginModule login is noninteractive or interactive.

- A noninteractive login does not prompt for any information whatsoever

- An interactive login prompts for principal name, password, or both

Noninteractive logins

The login proceeds noninteractively when you specify the credential type as initiator (`credsType=initiator`) and you perform one of the following actions:

- Specify the `useCcache` option
- Set the `useDefaultCcache` option to `true`

The login also proceeds noninteractively when you specify the credential type as acceptor or both (`credsType=acceptor` or `credsType=both`) and you perform one of the following actions:

- Specify the `useKeytab` option
- Set the `useDefaultKeytab` option to `true`

Interactive logins

Other configurations result in the login module prompting for a principal name and password so that it may obtain a TGT from a Kerberos KDC. The login module prompts for only a password when you specify the `principal` option.

Interactive logins require that the application specify `com.ibm.security.auth.callback.Krb5CallbackHandler` as the callback handler when creating the login context. The callback handler is responsible for prompting for input.

Credential type option

When you require the credential type to be both initiator and acceptor (`credsType=both`), `Krb5LoginModule` obtains both a TGT and a secret key. The login module uses the TGT to initiate contexts and the secret key to accept contexts. The JAAS configuration file must contain sufficient information to enable the login module to acquire the two types of credentials.

For credential types `acceptor` and `both`, the login module assumes a service principal. [«](#) [»](#)

Configuration and policy files

JGSS and JAAS depend on several configuration and policy files. You need to edit these files to conform to your environment and application. If you do not use JAAS with JGSS, you can safely ignore the JAAS configuration and policy files.

- Kerberos configuration file
- JAAS configuration file
- JAAS authorization policy file
- Java master security properties file
- Credentials cache and server key table

Note: In the following instructions, `{java.home}` denotes the path to the location of the version of Java that you are using on your server. For example, if you are using J2SDK, version 1.4, `{java.home}` is `/QIBM/ProdData/Java400/jdk14`. Remember to replace `{java.home}` in the property settings with the actual path to the Java home directory.

Kerberos configuration file

IBM JGSS requires a Kerberos configuration file. The default name and location of the Kerberos configuration file depends on the operating system being used. JGSS uses the following order to search for the default configuration file:

1. The file referenced by the Java property `java.security.krb5.conf`
2. `{java.home}/lib/security/krb5.conf`

3. c:\winnt\krb5.ini on Microsoft Windows^(R) platforms
4. /etc/krb5/krb5.conf on Solaris^(TM) platforms
5. /etc/krb5.conf on other Unix^(R) platforms

JAAS configuration file

The use of the JAAS login feature requires a JAAS configuration file. You can specify the JAAS configuration file by setting one of the following properties:

- The Java system property `java.security.auth.login.config`
- The security property `login.config.url.<integer>` in the `$(java.home)/lib/security/java.security` file

For more information, see the Sun Java Authentication and Authorization Service (JAAS) Web site .

JAAS policy file

When using the default policy implementation, JGSS grants JAAS permissions to entities by recording the permissions in a policy file. You can specify the JAAS policy file by setting one of the following properties:

- The Java system property `java.security.policy`
- The security property `policy.url.<integer>` in the `$(java.home)/lib/security/java.security` file

If you are using J2SDK, version 1.4, specifying a separate policy file for JAAS is optional. The default policy provider in J2SDK, version 1.4 supports the policy file entries that JAAS requires.

For more information, see the Sun Java Authentication and Authorization Service (JAAS) Web site .

Java master security properties file

A Java virtual machine (JVM) uses many important security properties that you set by editing the Java master security properties file. This file, named `java.security`, usually resides in the `$(java.home)/lib/security` directory on your iSeries server.

The following list describes several relevant security properties for using JGSS. Use the descriptions as a guide for editing the `java.security` file.

Note: When applicable, the descriptions include appropriate values required to run the JGSS samples.

security.provider.<integer>: The JGSS provider that you want to use. Also statically registers cryptographic provider classes. IBM JGSS uses cryptographic and other security services provided by the IBM JCE Provider. Specify the `sun.security.provider.Sun` and `com.ibm.crypto.provider.IBMJCE` packages exactly like the following example:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.crypto.provider.IBMJCE
```

policy.provider: System policy handler class. For example:

```
policy.provider=sun.security.provider.PolicyFile
```

policy.url.<integer>: URLs of policy files. To use the sample policy file, include an entry such as:

```
policy.url.1=file:/home/user/jgss/config/java.policy
```

login.configuration.provider: JAAS login configuration handler class, for example:

```
login.configuration.provider=com.ibm.security.auth.login.ConfigFile
```

auth.policy.provider: JAAS principal-based access control policy handler class, for example:

```
auth.policy.provider=com.ibm.security.auth.PolicyFile
```

login.config.url.<integer>: URLs for JAAS login configuration files. To use the sample configuration file, include an entry similar to:

```
login.config.url.1=file:/home/user/jgss/config/jaas.conf
```

auth.policy.url.<integer>: URLs for JAAS policy files. You can include both principal-based and CodeSource-based constructs in the JAAS policy file. To use the sample policy file, include an entry such as:

```
auth.policy.url.1=file:/home/user/jgss/config/jaas.policy
```

Credentials cache and server key table

A user principal keeps its Kerberos credentials in a credentials cache. A service principal keeps its secret key in a key table. At runtime, IBM JGSS locates these caches in the following ways:

User credentials cache

JGSS uses the following order to locate the user credentials cache:

1. The file referenced by the Java^(TM) property KRB5CCNAME
2. The file referenced by the environment variable KRB5CCNAME
3. /tmp/krb5cc_<uid> on Unix systems
4. \${user.home}/krb5cc_\${user.name}
5. \${user.home}/krb5cc (if \${user.name} cannot be obtained)

Server key table

JGSS uses the following order to locate the server key table file:

1. The value of the Java^(TM) property KRB5_KTNAME
2. default_keytab_name entry in the libdefaults stanza of the Kerberos configuration file
3. \${user.home}/krb5_keytab <<



Developing IBM JGSS applications

For more information about developing IBM JGSS applications, see the following topics:

Programming steps

Learn about the steps required to develop a JGSS application, including using transport tokens, creating the necessary JGSS objects, establishing and deleting context, and using per-message services.

Using JAAS with your JGSS application

Read about enabling the JAAS Kerberos login feature of JGSS. Information includes requirements for using the login feature and a snippet of example code.

Debugging

Read about using JGSS debugging to categorize and display helpful debugging messages.




JGSS javadoc reference information

Review the javadoc information for classes and methods in the org.ietf.jgss api package and for the Java versions of Kerberos credential management tools (kinit, ktab, and klist).

JGSS samples

Use the sample programs to discover how you can use JGSS in your applications. The sample documentation includes Java source code, instructions for running the samples, configuration and policy files, and more.

To develop JGSS applications, you need to be familiar with the high-level GSS-API specification and the Java bindings specification. IBM JGSS 1.0 is primarily based on and conforms to these specifications. See the following links for more information.

- RFC 2743: Generic Security Service Application Programming Interface Version 2, Update 1 
- RFC 2853: Generic Security Service API Version 2: Java Bindings  



IBM JGSS application programming steps

Operations in a JGSS application follow the Generic Security Service Application Programming Interface (GSS-API) operational model. For information about concepts important to JGSS operations, see JGSS concepts.


JGSS transport tokens

Some of the important JGSS operations generate tokens in the form of Java byte arrays. It is the responsibility of the application to forward the tokens from one JGSS peer to the other. JGSS does not constrain in any way the protocol that the application uses for transporting tokens. Applications may transport JGSS tokens together with other application (that is, non-JGSS) data. However, JGSS operations accept and use only JGSS-specific tokens.

Sequence of operations in a JGSS application

JGSS operations require certain programming constructs that you must use in the order listed below. Each of the steps applies to both the initiator and the acceptor.

Note: The information includes snippets of example code that illustrate using the high-level JGSS APIs and assume that your application imports the `org.ietf.jgss` package. Although many of the high-level APIs are overloaded, the snippets show only the most commonly used forms of those methods. Of course, use the API methods that best suit your needs.

1. Creating a GSSManager
An instance of GSSManager acts as a factory for creating other JGSS object instances.
2. Creating a GSSName
A GSSName represents the identity of a JGSS principal. Some JGSS operations can locate and use a default principal when you specify a null GSSName.
3. Creating a GSSCredential
A GSSCredential embodies the mechanism-specific credentials of the principal.
4. Creating a GSSContext
A GSSContext is used for context establishment and subsequent per-message services.
5. Selecting optional services on the context
Your application must explicitly request optional services, such as mutual authentication.
6. Establishing context
The initiator authenticates itself to the acceptor. However, when requesting mutual authentication, the acceptor in turn authenticates itself to the initiator.
7. Using per-message services
The initiator and the acceptor exchange secure messages over the established context.
8. Deleting context
The application deletes a context that is no longer needed. 



Creating a GSSManager

The GSSManager abstract class serves as a factory for creating the following JGSS objects:

- GSSName

- GSSCredential
- GSSContext

GSSManager also has methods for determining the supported security mechanisms and name types and for specifying JGSS providers. Use the GSSManager getInstance static method to create an instance of the default GSSManager:

```
GSSManager manager = GSSManager.getInstance();
```



Creating a GSSName

GSSName represents the identity of a GSS-API principal. A GSSName may contain many representations of the principal, one for each supported underlying mechanism. A GSSName that contains only one name representation is called a Mechanism Name (MN).

GSSManager has several overloaded methods for creating a GSSName from a string or a contiguous array of bytes. The methods interpret the string or byte array according to a specified name type. Typically, you use the GSSName byte-array methods to reconstitute an exported name. The exported name is usually a mechanism name of type GSSName.NT_EXPORT_NAME. Some of these methods allow you to specify a security mechanism with which to create the name.

Examples: Using GSSName

The following basic code snippet shows how to use GSSName.

Note: Specify Kerberos service name strings as either <service> or <service@host> where <service> is the name of the service and <host> is the hostname of the machine on which the service runs. You can (but do not have to) use a fully qualified hostname. When you omit the @<host> portion of the string, GSSName uses the local hostname.

```
// Create GSSName for user foo.
GSSName fooName = manager.createName("foo", GSSName.NT_USER_NAME);

// Create a Kerberos V5 mechanism name for user foo.
Oid krb5Mech = new Oid("1.2.840.113554.1.2.2");
GSSName fooName = manager.createName("foo", GSSName.NT_USER_NAME, krb5Mech);

// Create a mechanism name from a non-mechanism name by using the GSSName
// canonicalize method.
GSSName fooName = manager.createName("foo", GSSName.NT_USER_NAME);
GSSName fooKrb5Name = fooName.canonicalize(krb5Mech);
```



Creating a GSSCredential

A GSSCredential contains all the cryptographic information necessary to create a context on behalf of a principal and can contain credential information for multiple mechanisms.

GSSManager has three credential creation methods. Two of the methods take for parameters a GSSName, the lifetime of the credential, one or more mechanisms from which to get credentials, and the credential usage type. The third method takes only a usage type and uses the default values for other parameters. Specifying a null mechanism also uses the default mechanism. Specifying a null array of mechanisms causes the method to return credentials for the default set of mechanisms.

Note: Because IBM JGSS supports only the Kerberos V5 mechanism, that is the default mechanism.

Your application can create only one of the three credentials types (*initiate*, *accept*, or *initiate and accept*) at a time.

- A context initiator creates *initiate* credentials
- An acceptor creates *accept* credentials
- An acceptor that also behaves as an initiator creates *initiate and accept* credentials.

Examples: Obtaining credentials

The following example obtains the default credentials for an initiator:

```
GSSCredentials fooCreds = manager.createCredentials(GSSCredential.INITIATE)
```

The following example obtains Kerberos V5 credentials for the initiator foo that have the default validity period:

```
GSSCredential fooCreds = manager.createCredential(fooName, GSSCredential.DEFAULT_LIFETIME,
                                                    krb5Mech,GSSCredential.INITIATE);
```

The following example obtains an all-default acceptor credential:

```
GSSCredential serverCreds = manager.createCredential(null, GSSCredential.DEFAULT_LIFETIME,
                                                    (Oid)null, GSSCredential.ACCEPT);
```



Creating GSSContext

IBM JGSS supports two methods provided by GSSManager for creating a context:

- A method used by the context initiator
- A method used by the acceptor

Note: GSSManager provides a third method for creating a context that involves recreating previously exported contexts. However, because IBM JGSS Kerberos V5 mechanism does not support the use of exported contexts, IBM JGSS does not support this method.

Your application can not use an initiator context for context acceptance, nor can it use an acceptor context for context initiation. Both supported methods for creating a context require a credential as input. When the value of the credential is null, JGSS uses the default credential.

Examples: Using GSSContext

The following example creates a context with which the principal (foo) can initiate a context with the peer (superSecureServer) on the host (securityCentral). The example specifies the peer as superSecureServer@securityCentral. The created context is valid for the default period:

```
GSSName serverName = manager.createName("superSecureServer@securityCentral",
                                        GSSName.NT_HOSTBASED_SERVICE, krb5Mech);
GSSContext fooContext = manager.createContext(serverName, krb5Mech, fooCreds,
                                             GSSCredential.DEFAULT_LIFETIME);
```

The following example creates a context for superSecureServer in order to accept contexts initiated by any peer:

```
GSSContext serverAcceptorContext = manager.createContext(serverCreds);
```

Note that your application can create and simultaneously use both types of contexts.  

Requesting optional security services

Your application can request any of several optional security services. IBM JGSS supports the following optional services:

- Delegation

- Mutual authentication
- Replay detection
- Out-of-sequence detection
- Available per-message confidentiality
- Available per-message integrity

To request an optional service, your application must explicitly request it by using the appropriate request method on the context. Only an initiator can request these optional services. The initiator must make the request before context establishment begins.

For more information about optional services, see [Optional Service Support in Internet Engineering Task Force \(IETF\) RFC 2743 Generic Security Services Application Programming Interface Version 2, Update 1](#)



Example: Requesting optional services

In the following example, a context (fooContext) makes requests to enable mutual authentication and delegation services:

```
fooContext.requestMutualAuth(true);
fooContext.requestCredDeleg(true);
```



Establishing context

The two communicating peers must establish a security context over which they can use per-message services. The initiator calls `initSecContext()` on its context, which returns a token to the initiator application. The initiator application transports the context token to the acceptor application. The acceptor calls `acceptSecContext()` on its context, specifying the context token received from the initiator. Depending on the underlying mechanism and the optional services that the initiator selected, `acceptSecContext()` might produce a token that the acceptor application has to forward to the initiator application. The initiator application then uses the received token to call `initSecContext()` one more time.

An application can make multiple calls to `GSSContext.initSecContext()` and `GSSContext.acceptSecContext()`. An application can also exchange multiple tokens with a peer during context establishment. Hence, the typical method of establishing context uses a loop to call `GSSContext.initSecContext()` or `GSSContext.acceptSecContext()` until the applications establish context.

Example: Establishing context

The following example illustrates the initiator (foo) side of context establishment:

```
byte array[] inToken = null; // The input token is null for the first call
int inTokenLen = 0;

do {
    byte[] outToken = fooContext.initSecContext(inToken, 0, inTokenLen);

    if (outToken != null) {
        send(outToken); // transport token to acceptor
    }

    if( !fooContext.isEstablished()) {
        inToken = receive(); // receive token from acceptor
        inTokenLen = inToken.length;
    }
} while (!fooContext.isEstablished());
```

The following example illustrates the acceptor side of context establishment:

```
// The acceptor code for establishing context may be the following:
do {
    byte[] inToken = receive(); // receive token from initiator
    byte[] outToken =
        serverAcceptorContext.acceptSecContext(inToken, 0, inToken.length);

    if (outToken != null) {
        send(outToken); // transport token to initiator
    }
} while (!serverAcceptorContext.isEstablished());
```



Using per-message services

After establishing a security context, two communicating peers can exchange secure messages over the established context. Either peer can originate a secure message, regardless of whether it served as initiator or acceptor when establishing context. To make the message secure, IBM JGSS computes a cryptographic message integrity code (MIC) over the message. Optionally, IBM JGSS can have the Kerberos V5 mechanism encrypt the message to help ensure privacy.

Sending messages

IBM JGSS provides two sets of methods for securing messages: `wrap()` and `getMIC()`.

Using `wrap()`

The `wrap` method performs the following actions:

- Computes an MIC
- Encrypts the message (optional)
- Returns a token

The calling application uses the `MessageProp` class in conjunction with `GSSContext` to specify whether to apply encryption to the message.

The returned token contains both the MIC and text of the message. The text of the message is either ciphertext (for an encrypted message) or the original plaintext (for messages that are not encrypted).

Using `getMIC()`

The `getMIC` method performs the following actions but can not encrypt the message:

- Computes an MIC
- Returns a token

The returned token contains only the computed MIC and does not include the original message. So, in addition to transporting the MIC token to the peer, the peer must somehow be made aware of the original message so that it can verify the MIC.

Example: Using per-message services to send a message

The following example shows how one peer (foo) can wrap a message for delivery to another peer (superSecureServer):

```
byte[] message = "Ready to roll!".getBytes();
MessageProp mprop = new MessageProp(true); // foo wants the message encrypted
byte[] wrappedMessage =
    fooContext.wrap(message, 0, message.length, mprop);
```

```

send(wrappedMessage); // transfer the wrapped message to superSecureServer

// This is how superSecureServer may obtain a MIC for delivery to foo:
byte[] message = "You bet!".getBytes();
MessageProp mprop = null; // superSecureServer is content with
                          // the default quality of protection

byte[] mic =
    serverAcceptorContext.getMIC(message, 0, message.length, mprop);
send(mic);
// send the MIC to foo. foo also needs the original message to verify the MIC

```

Receiving messages

The receiver of a wrapped message uses `unwrap()` to decode the message. The `unwrap` method performs the following actions:

- Verifies the cryptographic MIC embedded in the message
- Returns the original message over which the sender computed the MIC

If the sender encrypted the message, `unwrap()` decrypts the message before verifying the MIC and then returns the original plaintext message. The receiver of an MIC token uses `verifyMIC()` to verify the MIC over a given a message.

The peer applications use their own protocol to deliver JGSS context and message tokens to each other. Peer applications also need to define a protocol for determining whether the token is an MIC or a wrapped message. For example, part of such a protocol may be as simple (and rigid) as that used by Simple Authentication and Security Layer (SASL) applications. The SASL protocol specifies that the context acceptor is always the first peer to send a per-message (wrapped) token following context establishment.

For more information, see Simple Authentication and Security Layer (SASL) .

Example: Using per-message services to receive a message

The following examples shows how a peer (`superSecureServer`) unwraps the wrapped token that it received from another peer (`foo`):

```

MessageProp mprop = new MessageProp(false);

byte[] plaintextFromFoo =
    serverAcceptorContext.unwrap(wrappedTokenFromFoo, 0,
                                wrappedTokenFromFoo.length, mprop);

// superSecureServer can now examine mprop to determine the message properties
// (such as whether the message was encrypted) applied by foo.

// foo verifies the MIC received from superSecureServer:

MessageProp mprop = new MessageProp(false);
fooContext.verifyMIC(micFromFoo, 0, micFromFoo.length, messageFromFoo, 0,
                    messageFromFoo.length, mprop);

// foo can now examine mprop to determine the message properties applied by
// superSecureServer. In particular, it can assert that the message was not
// encrypted since getMIC should never encrypt a message.

```



Deleting context

A peer deletes a context when the context is no longer needed. In JGSS operations, each peer unilaterally decides when to delete a context and does not need to inform its peer.

JGSS does not define a delete context token. To delete a context, the peer calls the dispose method of the GSSContext object to free up any resources used by the context. A disposed GSSContext object may still be accessible, unless the application sets the object to null. However, any attempt to use a disposed (but still accessible) context throws an exception. << >>

Using JAAS with your JGSS application

The IBM JGSS includes an optional JAAS login facility that allows the application to use JAAS to obtain credentials. After the JAAS login facility saves principal credentials and secret keys in the subject object of a JAAS login context, JGSS can retrieve the credentials from that subject.

The default behavior of JGSS is to retrieve credentials and secret keys from the subject. You can disable this feature by setting the Java property `javax.security.auth.useSubjectCredsOnly` to false.

Note: Although the pure Java JGSS provider can use the login interface, the native iSeries JGSS provider can not.

For more information about JAAS features, see [Obtaining Kerberos credentials and secret keys](#).

To use the JAAS login facility, your application must follow the JAAS programming model in the following ways:

- Create a JAAS login context
- Operate within the confines of a JAAS `Subject.doAs` construct

The following code snippet illustrates the concept of operating within the confines of a JAAS `Subject.doAs` construction:

```
static class JGSSOperations implements PrivilegedExceptionAction {
    public JGSSOperations() {}
    public Object run () throws GSSException {
        // JGSS application code goes/runs here
    }
}

public static void main(String args[]) throws Exception {
    // Create a login context that will use the Kerberos
    // callback handler
    // com.ibm.security.auth.callback.Krb5CallbackHandler

    // There must be a JAAS configuration for "JGSSClient"
    LoginContext loginContext =
        new LoginContext("JGSSClient", new Krb5CallabackHandler());
    loginContext.login();

    // Run the entire JGSS application in JAAS privileged mode
    Subject.doAsPrivileged(loginContext.getSubject(),
        new JGSSOperations(), null);
}
```



Debugging

When you are trying to identify JGSS problems, use the JGSS debugging capability to produce helpful categorized messages. You can turn on one or more categories by setting the appropriate values for the Java property `com.ibm.security.jgss.debug`. Activate multiple categories by using a comma to separate the category names.

Debugging categories include the following:

Category	Description
help	List debug categories
all	Turn on debugging for all categories
off	Turn off debugging completely
app	Application debugging (default)
ctx	Context operations debugging
cred	Credentials (including name) operations
marsh	Marshaling of tokens
mic	MIC operations
prov	Provider operations
qop	QOP operations
unmarsh	Unmarshaling of tokens
unwrap	Unwrap operations
wrap	Wrap operations

JGSS Debug class

To debug your JGSS application programmatically, use the debug class in the IBM JGSS framework. Your application can use the debug class to turn on and off debug categories and display debugging information for the active categories.

The default debugging constructor reads the Java property `com.ibm.security.jgss.debug` to determine which categories to activate (turn on).

Example: Debugging for the application category

The following example shows how to request debug information for the application category:

```
import com.ibm.security.jgss.debug;

Debug debug = new Debug(); // Gets categories from Java property

// Lots of work required to set up someBuffer. Test that the
// category is on before setting up for debugging.

if (debug.on(Debug.OPTS_CAT_APPLICATION)) {
    // Fill someBuffer with data.
    debug.out(Debug.OPTS_CAT_APPLICATION, someBuffer);
    // someBuffer may be a byte array or a String.
```



Samples: IBM Java Generic Security Service (JGSS)

The IBM Java Generic Security Service (JGSS) sample files include client and server programs, configuration files, policy files, and javadoc reference information.

You can view HTML versions of the samples or download the javadoc information and source code for the sample programs. Downloading the samples enables you to view the javadoc reference information, examine the code, edit the configuration and policy files, and compile and run the sample programs:

- View HTML versions of the samples

- Download and view the sample javadoc information
- Download and run the sample programs

Description of the sample programs

The JGSS samples include four programs:

- non-JAAS server
- non-JAAS client
- JAAS-enabled server
- JAAS-enabled client

The JAAS-enabled versions are fully interoperable with their non-JAAS counterparts. So, you can run a JAAS-enabled client against a non-JAAS server and you can run a non-JAAS client against a JAAS-enabled server.

Note: When you run a sample, you can specify one or more optional Java properties, including the names of the configuration and policy files, JGSS debug options, and the security manager. You can also turn on and turn off the JAAS features.

You can run the samples in either a one-server or a two-server configuration. The one server configuration consists of a client communicating with a primary server. The two-server configuration consists of a primary and a secondary server, where the primary server acts as an initiator, or client, to the secondary server.

When using a two-server configuration, the client first initiates a context and exchanges secure messages with the primary server. Next, the client delegates its credentials to the primary server. Then, on behalf of the client, the primary server uses these credentials to initiate a context and exchange secure messages with the secondary server. You can also use a two-server configuration in which the primary server acts as a client on its own behalf. In this case, the primary server uses its own credentials to initiate a context and exchange messages with the secondary server.

You can simultaneously run any number of clients against the primary server. Although you can run a client directly against the secondary server, the secondary server can not use delegated credentials or run as an initiator using its own credentials. << >>

Viewing the IBM JGSS samples

The IBM Java Generic Security Service (JGSS) sample files include client and server programs, configuration files, policy files, and javadoc reference information. Use the following links to view HTML versions of the JGSS samples.

For additional information, see the following topics:

- Description of the sample programs
- Downloading and running the sample programs

Viewing the sample programs

View the HTML versions of the JGSS sample programs by using the following links:

- Sample non-JAAS client program
- Sample non-JAAS server program
- Sample JAAS-enabled client program
- Sample JAAS-enabled server program

Viewing the sample configuration and policy files

View the HTML versions of the JGSS configuration and policy files by using the following links:

- Kerberos configuration file
- JAAS configuration file
- JAAS policy file
- Java policy file



Sample: IBM JGSS non-JAAS client program

For more information about using the sample client program, see [Downloading and running the sample programs](#).

Note: Read the Code example disclaimer for important legal information.

```
// IBM JGSS 1.0 Sample Client Program

package com.ibm.security.jgss.test;
import org.ietf.jgss.*;
import com.ibm.security.jgss.Debug;

import java.io.*;
import java.net.*;
import java.util.*;

/**
 * A JGSS sample client;
 * to be used in conjunction with the JGSS sample server.
 * The client first establishes a context with the server
 * and then sends wrapped message followed by a MIC to the server.
 * The MIC is calculated over the plain text that was wrapped.
 * The client requires to server to authenticate itself
 * (mutual authentication) during context establishment.
 * It also delegates its credentials to the server.
 *
 * It sets the JAVA variable
 * javax.security.auth.useSubjectCredsOnly to false
 * so that JGSS will not acquire credentials through JAAS.
 *
 * The client takes input parameters, and complements it
 * with information from the jgss.ini file; any required input not
 * supplied on the command line is taking from the jgss.ini file.
 *
 * Usage: Client [options]
 *
 * The -? option produces a help message including supported options.
 *
 * This sample client does not use JAAS.
 * The client can be run against the JAAS sample client and server.
 * See {@link JAASClient JAASClient} for a sample client that uses JAAS.
 */

class Client
{
    private Util testUtil      = null;
    private String myName     = null;
    private GSSName gssName   = null;
    private String serverName = null;
    private int servicePort   = 0;
    private GSSManager mgr    = GSSManager.getInstance();
    private GSSName service   = null;
    private GSSContext context = null;
    private String program    = "Client";
}
```



```

private String debugPrefix = "Client: ";
private TCPComms tcp      = null;
private String data       = null;
private byte[] dataBytes  = null;
private String serviceHostname = null;
private GSSCredential gssCred = null;

private static Debug debug      = new Debug();

private static final String usageString =
    "\t[-?] [-d | -n name] [-s serverName]"
    + "\n\t[-h serverHost [:port]] [-p port] [-m msg]"
    + "\n"
    + "\n  -?\t\t\tthe help; produces this message"
    + "\n  -n name\t\tthe client's principal name (without realm)"
    + "\n  -s serverName\t\tthe server's principal name (without realm)"
    + "\n  -h serverHost[:port]\tthe server's hostname"
    + "    " (and optional port number)"
    + "\n  -p port\t\tthe port on which the server will be listening"
    + "\n  -m msg\t\tmessage to send to the server";

// Caller must call initialize (may need to call processArgs first).
public Client (String programName) throws Exception
{
    testUtil = new Util();
    if (programName != null)
    {
        program = programName;
        debugPrefix = programName + ": ";
    }
}

// Caller must call initialize (may need to call processArgs first).
Client (String programName, boolean useSubjectCredsOnly) throws Exception
{
    this(programName);
    setUseSubjectCredsOnly(useSubjectCredsOnly);
}

public Client(GSSCredential myCred,
             String serverNameWithoutRealm,
             String serverHostname,
             int serverPort,
             String message)
    throws Exception
{
    testUtil = new Util();

    if (myCred != null)
    {
        gssCred = myCred;
    }
    else
    {
        throw new GSSEException(GSSEException.NO_CRED, 0,
                                "Null input credential");
    }

    init(serverNameWithoutRealm, serverHostname, serverPort, message);
}

void setUseSubjectCredsOnly(boolean useSubjectCredsOnly)
{
    final String subjectOnly = useSubjectCredsOnly ? "true" : "false";
    final String property = "javax.security.auth.useSubjectCredsOnly";
}

```

```

String temp = (String)java.security.AccessController.doPrivileged(
    new sun.security.action.GetPropertyAction(property));

if (temp == null)
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
        + "setting useSubjectCredsOnly property to "
        + useSubjectCredsOnly);

    // Property not set. Set it to the specified value.

    java.security.AccessController.doPrivileged(
        new java.security.PrivilegedAction() {
            public Object run() {
                System.setProperty(property, subjectOnly);
                return null;
            }
        }
    ));
}
else
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
        + "useSubjectCredsOnly property already set "
        + "in JVM to " + temp);
}
}

private void init(String myNameWithoutRealm,
    String serverNameWithoutRealm,
    String serverHostname,
    int serverPort,
    String message) throws Exception
{
    myName = myNameWithoutRealm;
    init(serverNameWithoutRealm, serverHostname, serverPort, message);
}

private void init(String serverNameWithoutRealm,
    String serverHostname,
    int serverPort,
    String message) throws Exception
{
    // peer's name
    if (serverNameWithoutRealm != null)
    {
        this.serverName = serverNameWithoutRealm;
    }
    else
    {
        this.serverName = testUtil.getDefaultServicePrincipalWithoutRealm();
    }

    // peer's host
    if (serverHostname != null)
    {
        this.serviceHostname = serverHostname;
    }
    else
    {
        this.serviceHostname = testUtil.getDefaultServiceHostname();
    }

    // peer's port
    if (serverPort > 0)
    {
        this.servicePort = serverPort;
    }
}

```

```

else
{
    this.servicePort = testUtil.getDefaultServicePort();
}

// message for peer
if (message != null)
{
    this.data = message;
}
else
{
    this.data = "The quick brown fox jumps over the lazy dog";
}

this.dataBytes = this.data.getBytes();

tcp = new TCPComms(serviceHostname, servicePort);
}

void initialize() throws Exception
{
    Oid krb5MechanismOid = new Oid("1.2.840.113554.1.2.2");

    if (gssCred == null)
    {
        if (myName != null)
        {
            debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                + "creating GSSName USER_NAME for "
                + myName);

            gssName = mgr.createName(
                myName,
                GSSName.NT_USER_NAME,
                krb5MechanismOid);

            debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                + "Canonicalized GSSName=" + gssName);
        }
        else
            gssName = null; // for default credentials

        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "creating"
            + ((gssName == null)? " default " : " ")
            + "credential");

        gssCred = mgr.createCredential(
            gssName,
            GSSCredential.DEFAULT_LIFETIME,
            (Oid)null,
            GSSCredential.INITIATE_ONLY);

        if (gssName == null)
        {
            gssName = gssCred.getName();

            myName = gssName.toString();

            debug.out(Debug.OPTS_CAT_APPLICATION,
                debugPrefix + "default credential principal=" + myName);
        }
    }

    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + gssCred);
}

```

```

debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
    + "creating canonicalized GSSName for serverName " + serverName);

service = mgr.createName(serverName,
    GSSName.NT_HOSTBASED_SERVICE,
    krb5MechanismOid);

debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
    + "Canonicalized server name = " + service);

debug.out(Debug.OPTS_CAT_APPLICATION,
    debugPrefix + "Raw data=" + data);
}

```

```

void establishContext(BitSet flags) throws Exception
{
    try {
        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "creating GSScontext");

        Oid defaultMech = null;
        context = mgr.createContext(service, defaultMech, gssCred,
            GSSContext.INDEFINITE_LIFETIME);

        if (flags != null)
        {
            if (flags.get(Util.CONTEXT_OPTS_MUTUAL))
            {
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                    + "requesting mutualAuthn");

                context.requestMutualAuth(true);
            }

            if (flags.get(Util.CONTEXT_OPTS_INTEG))
            {
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                    + "requesting integrity");

                context.requestInteg(true);
            }

            if (flags.get(Util.CONTEXT_OPTS_CONF))
            {
                context.requestConf(true);
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                    + "requesting confidentiality");
            }

            if (flags.get(Util.CONTEXT_OPTS_DELEG))
            {
                context.requestCredDeleg(true);
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                    + "requesting delegation");
            }

            if (flags.get(Util.CONTEXT_OPTS_REPLAY))
            {
                context.requestReplayDet(true);
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                    + "requesting replay detection");
            }

            if (flags.get(Util.CONTEXT_OPTS_SEQ))
            {

```

```

        context.requestSequenceDet(true);
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "requesting out-of-sequence detection");
    }
    // Add more later!
}

byte[] response = null;
byte[] request = null;
int len = 0;
boolean done = false;
do {
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
        + "Calling initSecContext");

    request = context.initSecContext(response, 0, len);

    if (request != null)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "Sending initial context token");

        tcp.send(request);
    }
    done = context.isEstablished();

    if (!done)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "Receiving response token");

        byte[] temp = tcp.receive();
        response = temp;
        len = response.length;
    }
} while(!done);

debug.out(Debug.OPTS_CAT_APPLICATION,
    debugPrefix + "context established with acceptor");

} catch (Exception exc) {
    exc.printStackTrace();
    throw exc;
}
}

void doMIC() throws Exception
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "generating MIC");
    byte[] mic = context.getMIC(dataBytes, 0, dataBytes.length, null);

    if (mic != null)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "sending MIC");
        tcp.send(mic);
    }
    else
        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "getMIC Failed");
}

void doWrap() throws Exception
{
    MessageProp mp = new MessageProp(true);
    mp.setPrivacy(context.getConfState());

    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "wrapping message");
}

```

```

byte[] wrapped = context.wrap(dataBytes, 0, dataBytes.length, mp);

if (wrapped != null)
{
    debug.out(Debug.OPTS_CAT_APPLICATION,
        debugPrefix + "sending wrapped message");

    tcp.send(wrapped);
}
else
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "wrap Failed");
}

void printUsage()
{
    System.out.println(program + usageString);
}

void processArgs(String[] args) throws Exception
{
    String port          = null;
    String myName        = null;
    int servicePort      = 0;
    String serviceHostname = null;

    String sHost = null;
    String msg = null;

    GetOptions options = new GetOptions(args, "?h:p:m:n:s:");
    int ch = -1;
    while ((ch = options.getopt()) != options.optEOF)
    {
        switch(ch)
        {
            case '?':
                printUsage();
                System.exit(1);

            case 'h':
                if (sHost == null)
                {
                    sHost = options.optArgGet();
                    int p = sHost.indexOf(':');
                    if (p != -1)
                    {
                        String temp1 = sHost.substring(0, p);
                        if (port == null)
                            port = sHost.substring(p+1, sHost.length()).trim();
                        sHost = temp1;
                    }
                }
                continue;

            case 'p':
                if (port == null)
                    port = options.optArgGet();
                continue;

            case 'm':
                if (msg == null)
                    msg = options.optArgGet();
                continue;

            case 'n':
                if (myName == null)
                    myName = options.optArgGet();

```

```

        continue;

        case 's':
            if (serverName == null)
                serverName = options.optArgGet();
            continue;
        }
    }

    if ((port != null) && (port.length() > 0))
    {
        int p = -1;
        try {
            p = Integer.parseInt(port);
        } catch (Exception exc) {
            System.out.println("Bad port input: "+port);
        }

        if (p != -1)
            servicePort = p;
    }

    if ((sHost != null) && (sHost.length() > 0)) {
        serviceHostname = sHost;
    }

    init(myName, serverName, serviceHostname, servicePort, msg);
}

void interactWithAcceptor(BitSet flags) throws Exception
{
    establishContext(flags);
    doWrap();
    doMIC();
}

void interactWithAcceptor() throws Exception
{
    BitSet flags = new BitSet();
    flags.set(Util.CONTEXT_OPTS_MUTUAL);
    flags.set(Util.CONTEXT_OPTS_CONF);
    flags.set(Util.CONTEXT_OPTS_INTEG);
    flags.set(Util.CONTEXT_OPTS_DELEG);
    interactWithAcceptor(flags);
}

void dispose() throws Exception
{
    if (tcp != null)
    {
        tcp.close();
    }
}

public static void main(String args[]) throws Exception
{
    System.out.println(debug.toString()); // XXXXXX
    String programName = "Client";
    Client client = null;
    try {
        client = new Client(programName,
                            false); // don't use Subject creds.
        client.processArgs(args);
        client.initialize();
        client.interactWithAcceptor();
    } catch (Exception exc) {
        debug.out(Debug.OPTS_CAT_APPLICATION,

```

```

                programName + " Exception: " + exc.toString());
        exc.printStackTrace();
        throw exc;
    } finally {
        try {
            if (client != null)
                client.dispose();
        } catch (Exception exc) {}
    }
}

debug.out(Debug.OPTS_CAT_APPLICATION, programName + ": done");
}
}

```



Sample: IBM JGSS non-JAAS server program

For more information about using the sample server program, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```

// IBM JGSS 1.0 Sample Server Program

package com.ibm.security.jgss.test;

import org.ietf.jgss.*;
import com.ibm.security.jgss.Debug;
import java.io.*;
import java.net.*;
import java.util.*;

/**
 * A JGSS sample server; to be used in conjunction with a JGSS sample client.
 *
 * It continuously listens for client connections,
 * spawning a thread to service an incoming connection.
 * It is capable of running multiple threads concurrently.
 * In other words, it can service multiple clients concurrently.
 *
 * Each thread first establishes a context with the client
 * and then waits for a wrapped message followed by a MIC.
 * It assumes that the client calculated the MIC over the plain
 * text wrapped by the client.
 *
 * If the client delegates its credential to the server, the delegated
 * credential is used to communicate with a secondary server.
 *
 * Also, the server can be started to act as a client as well as
 * a server (using the -b option). In this case, the first
 * thread spawned by the server uses the server principal's own credential
 * to communicate with the secondary server.
 *
 * The secondary server must have been started prior to the (primary) server
 * initiating contact with it (the secondary server).
 * In communicating with the secondary server, the primary server acts as
 * a JGSS initiator (i.e., client), establishing a context and engaging in
 * wrap and MIC per-message exchanges with the secondary server.
 *
 * The server takes input parameters, and complements it
 * with information from the jgss.ini file; any required input not
 * supplied on the command line is taking from the jgss.ini file.
 * Built-in defaults are used if there is no jgss.ini file or if a particular
 * variable is not specified in the ini file.
 *
 */

```



```

* Usage: Server [options]
*
* The -? option produces a help message including supported options.
*
* This sample server does not use JAAS.
* It sets the JAVA variable
* javax.security.auth.useSubjectCredsOnly to false
* so that JGSS will not acquire credentials through JAAS.
* The server can be run against the JAAS sample clients and servers.
* See {@link JAASServer JAASServer} for a sample server that uses JAAS.
*/

```

```

class Server implements Runnable
{

```

```

    /*
    * NOTES:
    * This class, Server, is expected to be run in concurrent
    * multiple threads. The static variables consist of variables
    * set from command-line arguments and variables (such as
    * the server's own credentials, gssCred) that are set once during
    * during initialization. These variables do not change
    * once set and are shared between all running threads.
    *
    * The only static variable that is changed after being set initially
    * is the variable 'beenInitiator' which is set 'true'
    * by the first thread to run the server as initiator using
    * the server's own creds. This ensures the server is run as an initiator
    * once only. Querying and modifying 'beenInitiator' is synchronized
    * between the threads.
    *
    * The variable 'tcp' is non-static and is set per thread
    * to represent the socket on which the client being serviced
    * by the thread connected.
    */

```

```

    private static Util testUtil          = null;
    private static int myPort             = 0;
    private static Debug debug            = new Debug();
    private static String myName          = null;
    private static GSSCredential gssCred  = null;
    private static String serviceNameNoRealm = null;
    private static String serviceHost     = null;
    private static int   servicePort     = 0;
    private static String serviceMsg      = null;
    private static GSSManager mgr         = null;
    private static GSSName gssName       = null;
    private static String program         = "Server";
    private static boolean clientServer   = false;
    private static boolean primaryServer  = true;

```

```

    private static boolean beenInitiator  = false;

```

```

    private static final String usageString =
        "\t[-?] [-# number] [-d | -n name] [-p port]"
        + "\n\t[-s serverName] [-h serverHost [:port]] [-P serverPort] [- msg]"
        + "\n"
        + "\n -?\t\tthe help; produces this message"
        + "\n -# number\t\tWhether primary or secondary server"
        + "      \n\t\t(1 = primary, 2 = secondary; default = first)"
        + "\n -n name\t\tthe server's principal name (without realm)"
        + "\n -p port\t\tthe port on which the server will be listening"
        + "\n -s serverName\tsecondary server's principal name"
        + "      (without realm)"
        + "\n -h serverHost[:port]\tsecondary server's hostname"
        + "      (and optional port number)"
        + "\n -P port\t\tsecondary server's port number"
        + "\n -m msg\t\tmessage to send to secondary server"

```

```

+ "\n -b \t\trun as both client and server"
+      " using the server's owns credentials";

// Non-static variables are thread-specific
// since each thread runs a separate instance of this class.

private String debugPrefix = null;
private TCPComms tcp      = null;

static {
    try {
        testUtil = new Util();
    } catch (Exception exc) {
        exc.printStackTrace();
        System.exit(1);
    }
}

Server (Socket socket) throws Exception
{
    debugPrefix = program + ": ";
    tcp = new TCPComms(socket);
}

Server (String program) throws Exception
{
    debugPrefix = program + ": ";
    this.program = program;
}

Server (String program, boolean useSubjectCredsOnly) throws Exception
{
    this(program);
    setUseSubjectCredsOnly(useSubjectCredsOnly);
}

void setUseSubjectCredsOnly(boolean useSubjectCredsOnly)
{
    final String subjectOnly = useSubjectCredsOnly ? "true" : "false";
    final String property = "javax.security.auth.useSubjectCredsOnly";

    String temp = (String)java.security.AccessController.doPrivileged(
        new sun.security.action.GetPropertyAction(property));

    if (temp == null)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "setting useSubjectCredsOnly property to "
            + (useSubjectCredsOnly ? "true" : "false"));

        // Property not set. Set it to the specified value.

        java.security.AccessController.doPrivileged(
            new java.security.PrivilegedAction() {
                public Object run() {
                    System.setProperty(property, subjectOnly);
                    return null;
                }
            });
    }
    else
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "useSubjectCredsOnly property already set "
            + "in JVM to " + temp);
    }
}

```

```

private void init(boolean primary,
                 String myNameWithoutRealm,
                 int port,
                 String serverNameWithoutRealm,
                 String serverHostname,
                 int serverPort,
                 String message,
                 boolean clientServer)
    throws Exception
{
    primaryServer = primary;
    this.clientServer = clientServer;

    myName = myNameWithoutRealm;

    // my port
    if (port > 0)
    {
        myPort = port;
    }
    else if (primary)
    {
        myPort = testUtil.getDefaultServicePort();
    }
    else
    {
        myPort = testUtil.getDefaultService2Port();
    }

    if (primary)
    {
        // peer's name
        if (serverNameWithoutRealm != null)
        {
            serviceNameNoRealm = serverNameWithoutRealm;
        }
        else
        {
            serviceNameNoRealm =
                testUtil.getDefaultService2PrincipalWithoutRealm();
        }

        // peer's host
        if (serverHostname != null)
        {
            if (serverHostname.equalsIgnoreCase("localhost"))
            {
                serverHostname = InetAddress.getLocalHost().getHostName();
            }

            serviceHost = serverHostname;
        }
        else
        {
            serviceHost = testUtil.getDefaultService2Hostname();
        }

        // peer's port
        if (serverPort > 0)
        {
            servicePort = serverPort;
        }
        else
        {
            servicePort = testUtil.getDefaultService2Port();
        }
    }
}

```

```

        // message for peer
        if (message != null)
        {
            serviceMsg = message;
        }
        else
        {
            serviceMsg = "Hi there! I am a server."
                + "But I can be a client, too";
        }
    }

    String temp = debugPrefix + "details"
        + "\n\tPrimary:\t" + primary
        + "\n\tName:\t\t" + myName
        + "\n\tPort:\t\t" + myPort
        + "\n\tClient+server:\t" + clientServer;

    if (primary)
    {
        temp += "\n\tOther Server:"
            + "\n\t\tName:\t" + serviceNameNoRealm
            + "\n\t\tHost:\t" + serviceHost
            + "\n\t\tPort:\t" + servicePort
            + "\n\t\tMsg:\t" + serviceMsg;
    }

    debug.out(Debug.OPTS_CAT_APPLICATION, temp);
}

void initialize() throws GSSException
{
    debug.out(Debug.OPTS_CAT_APPLICATION,
        debugPrefix + "creating GSSManager");

    mgr = GSSManager.getInstance();

    int usage = clientServer ? GSSCredential.INITIATE_AND_ACCEPT
        : GSSCredential.ACCEPT_ONLY;

    if (myName != null)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "creating GSSName for " + myName);

        gssName = mgr.createName(myName,
            GSSName.NT_HOSTBASED_SERVICE);

        Oid krb5MechanismOid = new Oid("1.2.840.113554.1.2.2");
        gssName.canonicalize(krb5MechanismOid);

        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "Canonicalized GSSName=" + gssName);
    }
    else
        gssName = null;

    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "creating"
        + ((gssName == null)? " default " : " ")
        + "credential");

    gssCred = mgr.createCredential(
        gssName, GSSCredential.DEFAULT_LIFETIME,
        (Oid)null, usage);

    if (gssName == null)
    {

```

```

        gssName = gssCred.getName();
        myName = gssName.toString();

        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "default credential principal=" + myName);
    }
}

```

```

void processArgs(String[] args) throws Exception
{
    String port    = null;
    String name    = null;
    int  iport    = 0;

    String sport   = null;
    int  isport   = 0;
    String sname   = null;
    String shost   = null;
    String smessage = null;

    boolean primary = true;
    String status   = null;

    boolean defaultPrinc = false;
    boolean clientServer = false;

    GetOptions options = new GetOptions(args, "?#:p:n:P:s:h:m:b");
    int ch = -1;
    while ((ch = options.getopt()) != options.optEOF)
    {
        switch(ch)
        {
            case '?':
                printUsage();
                System.exit(1);

            case '#':
                if (status == null)
                    status = options.optArgGet();
                continue;

            case 'p':
                if (port == null)
                    port = options.optArgGet();
                continue;

            case 'n':
                if (name == null)
                    name = options.optArgGet();
                continue;

            case 'b':
                clientServer = true;
                continue;

            ////// The other server

            case 'P':
                if (sport == null)
                    sport = options.optArgGet();
                continue;

            case 'm':
                if (smessage == null)
                    smessage = options.optArgGet();

```

```

        continue;

    case 's':
        if (sname == null)
            sname = options.optArgGet();
        continue;

    case 'h':
        if (shost == null)
        {
            shost = options.optArgGet();
            int p = shost.indexOf(':');
            if (p != -1)
            {
                String temp1 = shost.substring(0, p);
                if (sport == null)
                    sport = shost.substring
                        (p+1, shost.length()).trim();
                shost = temp1;
            }
        }
        continue;
    }
}

if (defaultPrinc && (name != null))
{
    System.out.println(
        "ERROR: '-d' and '-n ' options are mutually exclusive");
    printUsage();
    System.exit(1);
}

if (status != null)
{
    int p = -1;
    try {
        p = Integer.parseInt(status);
    } catch (Exception exc) {
        System.out.println( "Bad status input: "+status);
    }

    if (p != -1)
    {
        primary = (p == 1);
    }
}

if (port != null)
{
    int p = -1;
    try {
        p = Integer.parseInt(port);
    } catch (Exception exc) {
        System.out.println( "Bad port input: "+port);
    }
    if (p != -1)
        iport = p;
}

if (sport != null)
{
    int p = -1;
    try {
        p = Integer.parseInt(sport);
    } catch (Exception exc) {
        System.out.println( "Bad server port input: "+port);
    }
}

```

```

    }
    if (p != -1)
        isport = p;
}

init(primary, // first or second server
     name,    // my name
     ipport,  // my port
     sname,   // other server's name
     shost,   // other server's hostname
     isport,  // other server's port
     smessage, // msg for other server
     clientServer); // whether to run as initiator with own creds
}

void processRequests() throws Exception
{
    ServerSocket ssocket = null;
    Server server = null;
    try {
        ssocket = new ServerSocket(myPort);
        do {
            debug.out(Debug.OPTS_CAT_APPLICATION,
                      debugPrefix + "listening on port " + myPort + " ...");
            Socket csocket = ssocket.accept();

            debug.out(Debug.OPTS_CAT_APPLICATION,
                      debugPrefix + "incoming connection on " + csocket);

            server = new Server(csocket); // set client socket per thread
            Thread thread = new Thread(server);
            thread.start();
            if (!thread.isAlive())
                server.dispose(); // close the client socket
        } while(true);
    } catch (Exception exc) {
        debug.out(Debug.OPTS_CAT_APPLICATION,
                  debugPrefix + "*** ERROR processing requests ***");
        exc.printStackTrace();
    } finally {
        try {
            if (ssocket != null)
                ssocket.close(); // close the server socket
            if (server != null)
                server.dispose(); // close the client socket
        } catch (Exception exc) {}
    }
}

void dispose()
{
    try {
        if (tcp != null)
        {
            tcp.close();
            tcp = null;
        }
    } catch (Exception exc) {}
}

boolean establishContext(GSSContext context) throws Exception
{
    byte[] response = null;
    byte[] request = null;

    debug.out(Debug.OPTS_CAT_APPLICATION,
              debugPrefix + "establishing context");
}

```

```

do {
    request = tcp.receive();
    if (request == null || request.length == 0)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "Received no data; perhaps client disconnected");

        return false;
    }

    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "accepting");
    if ((response = context.acceptSecContext
        (request, 0, request.length)) != null)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "sending response");
        tcp.send(response);
    }
} while(!context.isEstablished());

debug.out(Debug.OPTS_CAT_APPLICATION,
    debugPrefix + "context established - " + context);

return true;
}

byte[] unwrap(GSSContext context, byte[] msg) throws Exception
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "unwrapping");

    MessageProp mp = new MessageProp(true);
    byte[] unwrappedMsg = context.unwrap(msg, 0, msg.length, mp);

    debug.out(Debug.OPTS_CAT_APPLICATION,
        debugPrefix + "unwrapped msg is:");
    debug.out(Debug.OPTS_CAT_APPLICATION, unwrappedMsg);

    return unwrappedMsg;
}

void verifyMIC (GSSContext context, byte[] mic, byte[] raw) throws Exception
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "verifying MIC");

    MessageProp mp = new MessageProp(true);
    context.verifyMIC(mic, 0, mic.length, raw, 0, raw.length, mp);

    debug.out(Debug.OPTS_CAT_APPLICATION,
        debugPrefix + "successfully verified MIC");
}

void useDelegatedCred(GSSContext context) throws Exception
{
    GSSCredential delCred = context.getDelegCred();
    if (delCred != null)
    {
        if (primaryServer)
        {
            debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix +
                "Primary server received delegated cred; using it");
            runAsInitiator(delCred); // using delegated creds
        }
        else
        {
            debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix +
                "Non-primary server received delegated cred; "

```



```

        + "ignoring it");
    }
}
else
{
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix +
              "ERROR: null delegated cred");
}
}

public void run()
{
    byte[] response        = null;
    byte[] request         = null;
    boolean unwrapped     = false;
    GSSContext context     = null;

    try {
        Thread currentThread = Thread.currentThread();
        String threadName    = currentThread.getName();

        debugPrefix = program + " " + threadName + ": ";

        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                  + "servicing client ...");

        debug.out(Debug.OPTS_CAT_APPLICATION,
                  debugPrefix + "creating GSSContext");

        context = mgr.createContext(gssCred);

        // First establish context with the initiator.
        if (!establishContext(context))
            return;

        // Then process messages from the initiator.
        // We expect to receive a wrapped message followed by a MIC.
        // The MIC should have been calculated over the plain
        // text that we received wrapped.
        // Use delegated creds if any.
        // Then run as initiator using own creds if necessary; only
        // the first thread does this.

        do {
            debug.out(Debug.OPTS_CAT_APPLICATION,
                      debugPrefix + "receiving per-message request");

            request = tcp.receive();
            if (request == null || request.length == 0)
            {
                debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
                          + "Received no data; perhaps client disconnected");

                return;
            }

            // Expect wrapped message first.
            if (!unwrapped)
            {
                response = unwrap(context, request);
                unwrapped = true;
                continue; // get next request
            }

            // Followed by a MIC.
            verifyMIC(context, request, response);
        }
    }
}

```

```

        // Impersonate the initiator if it delegated its creds to us.
        if (context.getCredDelegState())
            useDelegatedCred(context);

        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "clientServer=" + clientServer
            + ", beenInitiator=" + beenInitiator);

        // If necessary, run as initiator using our own creds.
        if (clientServer)
            runAsInitiatorOnce(currentThread);

        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "done");
        return;
    } while(true);
} catch (Exception exc) {
    debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix + "ERROR");
    exc.printStackTrace();

    // Squelch per-thread exceptions so we don't bring
    // the server down because of exceptions in
    // individual threads.
    return;
} finally {
    if (context != null)
    {
        try {
            context.dispose();
        } catch (Exception exc) {}
    }
}
}

synchronized void runAsInitiatorOnce(Thread thread)
    throws InterruptedException
{
    if (!beenInitiator)
    {
        // set flag true early to prevent subsequent threads
        // from attempting to runAsInitiator.
        beenInitiator = true;

        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix +
            "About to run as initiator with own creds ...");

        //thread.sleep(30*1000, 0);
        runAsInitiator();
    }
}

void runAsInitiator(GSSCredential cred)
{
    Client client = null;
    try {
        client = new Client(cred,
            serviceNameNoRealm,
            serviceHost,
            servicePort,
            serviceMsg);

        client.initialize();

        BitSet flags = new BitSet();

```

```

        flags.set(Util.CONTEXT_OPTS_MUTUAL);
        flags.set(Util.CONTEXT_OPTS_CONF);
        flags.set(Util.CONTEXT_OPTS_INTEG);

        client.interactWithAcceptor(flags);
    } catch (Exception exc) {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "Exception running as initiator");

        exc.printStackTrace();
    } finally {
        try {
            client.dispose();
        } catch (Exception exc) {}
    }
}

void runAsInitiator()
{
    if (clientServer)
    {
        debug.out(Debug.OPTS_CAT_APPLICATION,
            debugPrefix + "running as initiator with own creds");

        runAsInitiator(gssCred); // use own creds;
    }
    else
    {
        debug.out(Debug.OPTS_CAT_APPLICATION, debugPrefix
            + "Cannot run as initiator with own creds "
            + "\nbecause not running as both initiator and acceptor.");
    }
}

void printUsage()
{
    System.out.println(program + usageString);
}

public static void main(String[] args) throws Exception
{
    System.out.println(debug.toString()); // XXXXXXXX
    String programName = "Server";
    try {
        Server server = new Server(programName,
            false); // don't use creds from Subject
        server.processArgs(args);
        server.initialize();
        server.processRequests();
    } catch (Exception exc) {
        debug.out(Debug.OPTS_CAT_APPLICATION, programName + ": EXCEPTION");
        exc.printStackTrace();
        throw exc;
    }
}
}

```



Sample: IBM JGSS JAAS-enabled client program

For more information about using the sample client program, see Downloading and running the IBM JGSS samples.

Note: Read the Code example disclaimer for important legal information.

```

// IBM Java GSS 1.0 sample JAAS-enabled client program

package com.ibm.security.jgss.test;
import com.ibm.security.jgss.Debug;
import com.ibm.security.auth.callback.Krb5CallbackHandler;
import javax.security.auth.Subject;
import javax.security.auth.login.LoginContext;
import java.security.PrivilegedExceptionAction;

/**
 * A Java GSS sample client that uses JAAS.
 *
 * It does a JAAS login and operates within the JAAS login context so created.
 *
 * It does not set the JAVA variable
 * javax.security.auth.useSubjectCredsOnly, leaving
 * the variable to default to true
 * so that Java GSS acquires credentials from the JAAS Subject
 * associated with login context (created by the client).
 *
 * The JAASClient is equivalent to its superclass {@link Client Client}
 * in all other respects, and it
 * can be run against the non-JAAS sample clients and servers.
 */

class JAASClient extends Client
{
    JAASClient(String programName) throws Exception
    {
        // Do not set useSubjectCredsOnly. Set only the program name.
        // useSubjectCredsOnly default to "true" if not set.
        super(programName);
    }

    static class JAASClientAction implements PrivilegedExceptionAction
    {
        private JAASClient client;

        public JAASClientAction(JAASClient client)
        {
            this.client = client;
        }

        public Object run () throws Exception
        {
            client.initialize();
            client.interactWithAcceptor();
            return null;
        }
    }

    public static void main(String args[]) throws Exception
    {
        String programName = "JAASClient";
        JAASClient client = null;
        Debug dbg = new Debug();

        System.out.println(dbg.toString()); // XXXXXXXX

        try {
            client = new JAASClient(programName); //use Subject creds
            client.processArgs(args);

            LoginContext loginCtx = new LoginContext("JAASClient",
                new Krb5CallbackHandler());

            loginCtx.login();
        }
    }
}

```

```

        dbg.out(Debug.OPTS_CAT_APPLICATION,
                programName + ": Kerberos login OK");

        Subject subject = loginCtxt.getSubject();

        PrivilegedExceptionAction jaasClientAction
            = new JAASClientAction(client);

        Subject.doAsPrivileged(subject, jaasClientAction, null);

    } catch (Exception exc) {
        dbg.out(Debug.OPTS_CAT_APPLICATION,
                programName + " Exception: " + exc.toString());
        exc.printStackTrace();
        throw exc;
    } finally {
        try {
            if (client != null)
                client.dispose();
        } catch (Exception exc) {}
    }

    dbg.out(Debug.OPTS_CAT_APPLICATION,
            programName + ": Done ...");
}
}

```



Sample: IBM JGSS JAAS-enabled server program

For more information about using the sample server program, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```

// IBM Java GSS 1.0 sample JAAS-enabled server program

package com.ibm.security.jgss.test;
import com.ibm.security.jgss.Debug;
import com.ibm.security.auth.callback.Krb5CallbackHandler;
import javax.security.auth.Subject;
import javax.security.auth.login.LoginContext;
import java.security.PrivilegedExceptionAction;

/**
 * A Java GSS sample server that uses JAAS.
 *
 * It does a JAAS login and operates within the JAAS login context so created.
 *
 * It does not set the JAVA variable
 * javax.security.auth.useSubjectCredsOnly, leaving
 * the variable to default to true
 * so that Java GSS acquires credentials from the JAAS Subject
 * associated with login context (created by the server).
 *
 * The JAASServer is equivalent to its superclass {@link Server Server}
 * in all other respects, and it
 * can be run against the non-JAAS sample clients and servers.
 */

class JAASServer extends Server
{
    JAASServer(String programName) throws Exception
    {

```

```

        super(programName);
    }

    static class JAASServerAction implements PrivilegedExceptionAction
    {
        private JAASServer server = null;

        JAASServerAction(JAASServer server)
        {
            this.server = server;
        }

        public Object run() throws Exception
        {
            server.initialize();
            server.processRequests();

            return null;
        }
    }

    public static void main(String[] args) throws Exception
    {
        String programName    = "JAASServer";
        Debug dbg              = new Debug();

        System.out.println(dbg.toString()); // XXXXXXXX

        try {
            // Do not set useSubjectCredsOnly.
            // useSubjectCredsOnly defaults to "true" if not set.

            JAASServer server = new JAASServer(programName);

            server.processArgs(args);

            LoginContext loginCtxt = new LoginContext(programName,
                new Krb5CallbackHandler());

            dbg.out(Debug.OPTS_CAT_APPLICATION, programName + ": Login in ...");

            loginCtxt.login();

            dbg.out(Debug.OPTS_CAT_APPLICATION, programName +
                ": Login successful");

            Subject subject = loginCtxt.getSubject();

            JAASServerAction serverAction = new JAASServerAction(server);

            Subject.doAsPrivileged(subject, serverAction, null);
        } catch (Exception exc) {
            dbg.out(Debug.OPTS_CAT_APPLICATION, programName + " EXCEPTION");
            exc.printStackTrace();
            throw exc;
        }
    }
}

```



Sample: Kerberos configuration file

For more information about using the sample configuration file, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```
# -----  
# Kerberos configuration file for running the JGSS sample applications.  
# Modify the entries to suit your environment.  
#-----  
  
[libdefaults]  
default_keytab_name = /QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab  
default_realm       = REALM.IBM.COM  
default_tkt_enctypes = des-cbc-crc  
default_tgs_enctypes = des-cbc-crc  
default_checksum    = rsa-md5  
kdc_timesync        = 0  
kdc_default_options = 0x40000010  
clockskew           = 300  
check_delegate      = 1  
ccache_type         = 3  
kdc_timeout         = 60000  
  
[realms]  
REALM.IBM.COM = {  
    kdc = kdc.ibm.com:88  
}  
  
[domain_realm]  
.ibm.com = REALM.IBM.COM
```



Sample: JAAS login configuration file

For more information about using the sample configuration file, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```
/**  
 * -----  
 * JAAS Login Configuration for the JGSS samples.  
 * -----  
 *  
 * Code example disclaimer  
 * IBM grants you a nonexclusive copyright license to use all programming code  
 * examples from which you can generate similar function tailored to your own  
 * specific needs.  
 * All sample code is provided by IBM for illustrative purposes only.  
 * These examples have not been thoroughly tested under all conditions.  
 * IBM, therefore, cannot guarantee or imply reliability, serviceability, or  
 * function of these programs.  
 * All programs contained herein are provided to you "AS IS" without any  
 * warranties of any kind.  
 * The implied warranties of non-infringement, merchantability and fitness  
 * for a particular purpose are expressly disclaimed.  
 *  
 *  
 * Supported options:  
 *   principal=<string>  
 *   credsType=initiator|acceptor|both (default=initiator)  
 *   forwardable=true|false (default=false)  
 *   proxiabile=true|false (default=false)  
 *   useCcache=<URL_string>  
 *   useKeytab=<URL_string>  
 *   useDefaultCcache=true|false (default=false)  
 *   useDefaultKeytab=true|false (default=false)  
 *   noAddress=true|false (default=false)
```

```

*
* Default realm (which is obtained from the Kerberos config file) is
* used if the principal specified does not include a realm component.
*/

```

```

JAASClient {
  com.ibm.security.auth.module.Krb5LoginModule required
  useDefaultCcache=true;
};

```

```

JAASServer {
  com.ibm.security.auth.module.Krb5LoginModule required
  credsType=acceptor useDefaultKeytab=true
  principal=gss_service/myhost.ibm.com@REALM.IBM.COM;
};

```



Sample: JAAS policy file

For more information about using the sample policy file, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```

// -----
// JAAS policy file for running the JGSS sample applications.
// Modify these permissions to suit your environment.
// Not recommended for use for any purpose other than that stated above.
// In particular, do not use this policy file or its
// contents to protect resources in a production environment.
//
// Code example disclaimer
// IBM grants you a nonexclusive copyright license to use all programming code
// examples from which you can generate similar function tailored to your own
// specific needs.
// All sample code is provided by IBM for illustrative purposes only.
// These examples have not been thoroughly tested under all conditions.
// IBM, therefore, cannot guarantee or imply reliability, serviceability, or
// function of these programs.
// All programs contained herein are provided to you "AS IS" without any
// warranties of any kind.
// The implied warranties of non-infringement, merchantability and fitness
// for a particular purpose are expressly disclaimed.
//
// -----
//-----
// Permissions for client only
//-----

grant CodeBase "file:ibmjgsssample.jar",
  Principal javax.security.auth.kerberos.KerberosPrincipal
    "bob@REALM.IBM.COM"
{
  // foo needs to be able to initiate a context with the server
  permission javax.security.auth.kerberos.ServicePermission
    "gss_service/myhost.ibm.com@REALM.IBM.COM", "initiate";

  // So that foo can delegate his creds to the server
  permission javax.security.auth.kerberos.DelegationPermission
    "\"gss_service/myhost.ibm.com@REALM.IBM.COM\" \"krbtgt/REALM.IBM.COM@REALM.IBM.COM\"";
};

//-----
// Permissions for the server only

```



```
//-----
grant CodeBase "file:ibmjgsssample.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "gss_service/myhost.ibm.com@REALM.IBM.COM"
{
    // Permission for the server to accept network connections on its host
    permission java.net.SocketPermission "myhost.ibm.com", "accept";

    // Permission for the server to accept JGSS contexts
    permission javax.security.auth.kerberos.ServicePermission
        "gss_service/myhost.ibm.com@REALM.IBM.COM", "accept";

    // The server acts as a client when communicating with the secondary (backup) server
    // This permission allows the server to initiate a context with the secondary server
    permission javax.security.auth.kerberos.ServicePermission
        "gss_service2/myhost.ibm.com@REALM.IBM.COM", "initiate";
};

//-----
// Permissions for the secondary server
//-----

grant CodeBase "file:ibmjgsssample.jar",
    Principal javax.security.auth.kerberos.KerberosPrincipal
        "gss_service2/myhost.ibm.com@REALM.IBM.COM"
{
    // Permission for the secondary server to accept network connections on its host
    permission java.net.SocketPermission "myhost.ibm.com", "accept";

    // Permission for the server to accept JGSS contexts
    permission javax.security.auth.kerberos.ServicePermission
        "gss_service2/myhost.ibm.com@REALM.IBM.COM", "accept";
};
```



Sample: Java policy file

For more information about using the sample policy file, see [Downloading and running the IBM JGSS samples](#).

Note: Read the Code example disclaimer for important legal information.

```
// -----
// Java policy file for running the JGSS sample applications on
// the iSeries server.
// Modify these permissions to suit your environment.
// Not recommended for use for any purpose other than that stated above.
// In particular, do not use this policy file or its
// contents to protect resources in a production environment.
//
// Code example disclaimer
// IBM grants you a nonexclusive copyright license to use all programming code
// examples from which you can generate similar function tailored to your own
// specific needs.
// All sample code is provided by IBM for illustrative purposes only.
// These examples have not been thoroughly tested under all conditions.
// IBM, therefore, cannot guarantee or imply reliability, serviceability, or
// function of these programs.
// All programs contained herein are provided to you "AS IS" without any
// warranties of any kind.
// The implied warranties of non-infringement, merchantability and fitness
// for a particular purpose are expressly disclaimed.
//
//-----
```

```

grant CodeBase "file:ibmjgsssample.jar" {
    // For Java 1.3
    permission javax.security.auth.AuthPermission "createLoginContext";

    // For Java 1.4
    permission javax.security.auth.AuthPermission "createLoginContext.JAASClient";
    permission javax.security.auth.AuthPermission "createLoginContext.JAASServer";

    permission javax.security.auth.AuthPermission "doAsPrivileged";

    // Permission to request a ticket from the KDC
    permission javax.security.auth.kerberos.ServicePermission
        "krbtgt/REALM.IBM.COM@REALM.IBM.COM", "initiate";

    // Permission to access sun.security.action classes
    permission java.lang.RuntimePermission "accessClassInPackage.sun.security.action";

    // A whole bunch of Java properties are accessed
    permission java.util.PropertyPermission "java.net.preferIPv4Stack", "read";
    permission java.util.PropertyPermission "java.version", "read";
    permission java.util.PropertyPermission "java.home", "read";
    permission java.util.PropertyPermission "user.home", "read";
    permission java.util.PropertyPermission "DEBUG", "read";
    permission java.util.PropertyPermission "com.ibm.security.jgss.debug", "read";
    permission java.util.PropertyPermission "java.security.krb5.kdc", "read";
    permission java.util.PropertyPermission "java.security.krb5.realm", "read";
    permission java.util.PropertyPermission "java.security.krb5.conf", "read";
    permission java.util.PropertyPermission "javax.security.auth.useSubjectCredsOnly", "read,write";

    // Permission to communicate with the Kerberos KDC host
    permission java.net.SocketPermission "kdc.ibm.com", "connect,accept,resolve";

    // I run the samples from my localhost
    permission java.net.SocketPermission "myhost.ibm.com", "accept,connect,resolve";
    permission java.net.SocketPermission "localhost", "listen,accept,connect,resolve";

    // Access to some possible Kerberos config locations
    // Modify the file paths as applicable to your environment
    permission java.io.FilePermission "${user.home}/krb5.ini", "read";
    permission java.io.FilePermission "${java.home}/lib/security/krb5.conf", "read";

    // Access to the Kerberos key table so we can get our server key.
    permission java.io.FilePermission "/QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab", "read";

    // Access to the user's Kerberos credentials cache.
    permission java.io.FilePermission "${user.home}/krb5cc_${user.name}", "read";
};

```



Samples: Downloading and viewing javadoc information for the IBM JGSS samples

To download and view the documentation for the IBM JGSS sample programs, complete the following steps:

1. Choose an existing directory (or create a new one) where you want to store the javadoc information.
2. Download the javadoc information (jgsssampledoc.zip) into the directory.
3. Extract the files from jgsssampledoc.zip into the directory.
4. Use your browser to access the index.htm file.

Code example disclaimer

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

All sample code is provided by IBM for illustrative purposes only. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

All programs contained herein are provided to you "AS IS" without any warranties of any kind. The implied warranties of non-infringement, merchantability and fitness for a particular purpose are expressly disclaimed. <>

Samples: Downloading and running the sample programs

Before modifying or running the samples, read the description of the sample programs.

To run the sample programs, perform the following tasks:

1. Download the sample files to your iSeries server
2. Prepare to run the sample files
3. Run the sample programs

For more information about how to run a sample, see Example: Running the non-JAAS sample.

Code example disclaimer

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

All sample code is provided by IBM for illustrative purposes only. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

All programs contained herein are provided to you "AS IS" without any warranties of any kind. The implied warranties of non-infringement, merchantability and fitness for a particular purpose are expressly disclaimed. <>

Samples: Downloading the IBM JGSS samples

Before modifying or running the samples, read the description of the sample programs.

To download the sample files and store them on your iSeries server, complete the following steps:

1. On your iSeries server, choose an existing directory (or create a new one) where you want to store the sample programs, configuration files, and policy files.
2. Download the sample programs (ibmjgsssample.zip).
3. Extract the files from ibmjgsssample.zip into the directory on the server.

Extracting the contents of ibmjgsssample.jar performs the following actions:

- Places ibmjgsssample.jar, which contains the sample .class files, into the selected directory
- Creates a subdirectory (named config) that contains the configuration and policy files
- Creates a subdirectory (named src) that contains the sample .java source files

Related information

You may want to read about related tasks or look at an example:

- Prepare to run the sample files
- Run the sample programs

- Example: Running the non-JAAS sample 



Samples: Preparing to run the sample programs


Before modifying or running the samples, read the description of the sample programs.

After you download the source code, you need to perform the following tasks before you can run the sample programs:

- Edit the configuration and policy files to suit your environment. For more information, refer to the comments in each configuration and policy file.
- Ensure that the `java.security` file contains the correct settings for your iSeries server. For more information, see Java master security properties file.
- Place the modified Kerberos configuration file (`krb5.conf`) into the directory on your iSeries server that is appropriate for the version of the J2SDK that you are using:
 - For Version 1.3 of the J2SDK: `/QIBM/ProdData/Java400/jdk13/lib/security`
 - For Version 1.4 of the J2SDK: `/QIBM/ProdData/Java400/jdk14/lib/security`

Related information

You may want to read about related tasks or look at an example::

- Download the sample files to your iSeries server
- Run the sample programs
- Example: Running the non-JAAS sample 



Samples: Running the sample programs

Before modifying or running the samples, read the description of the sample programs.

After you download and modify the source code, you can run one of the samples.

To run a sample, you must start the server program first. The server program must be running and ready to receive connections before you start the client program. The server is ready to receive connections when you see `listening on port <server_port>`. Make sure to remember or write down the `<server_port >`, which is the port number that you need to specify when you start the client.

Use the following command to start a sample program:

```
java [-Dproperty1=value1 ... -DpropertyN=valueN] com.ibm.security.jgss.test.<program> [options]
```

where

- `[-DpropertyN=valueN]` is one or more optional Java properties, including the names of the configuration and policy files, JGSS debug options, and the security manager. For more information, see the following example and Running JGSS applications.
- `<program>` is a required parameter that specifies the sample program that you want to run (either `Client`, `Server`, `JAASClient`, or `JAASServer`).
- `[options]` is an optional parameter for the sample program that you want to run. To display a list of supported options, use the following command:


```
java com.ibm.security.jgss.test.<program> -?
```

Note: Turn off the JAAS features in a JGSS-enabled sample by setting the Java property `javax.security.auth.useSubjectCredsOnly` to `false`. Of course, the default value of the JAAS-enabled

samples is to turn on JAAS, meaning that the property value is true. The non-JAAS client and server programs set the property to false, unless you have explicitly set the property value.

Related information

You may want to read about related tasks or look at an example:


- Download the sample files to your iSeries server
- Prepare to run the sample files
- Example: Running the non-JAAS sample 



IBM JGSS javadoc reference information

The javadoc reference information for IBM JGSS includes classes and methods in the org.ietf.jgss api package and the Java versions of some Kerberos credential management tools.

Although JGSS includes several publicly accessible packages (for example, com.ibm.security.jgss and com.ibm.security.jgss.spi), you should use only APIs from the standardized org.ietf.jgss package. Using only this package ensures that your application conforms to the GSS-API specifications and ensures optimum interoperability and portability.

- org.ietf.jgss
- kinit
- ktab
- klist 

Chapter 5. Tune Java program performance with the IBM Developer Kit for Java

You should take several aspects of Java^(TM) application performance into consideration when building a Java application for your iSeries server. Here are some links to details and hints on how you can get better performance:

- Use the Create Java Program (CRTJVAPGM) command before running Java class files, JAR files, or ZIP files to improve the runtime performance of your Java code.
- Change your optimization levels to achieve the best static compilation performance.
- Carefully set your values for optimal garbage collection performance.
- Only use native methods to start system functions that are relatively long running and are not available directly in Java.
- Use the `javac -o` option at compilation time to perform method inlining and significantly improve your method call performance.
- Use Java exceptions in cases that are not the normal flow through your application.

Use these tools with the Performance Explorer (PEX) to locate performance problems in your Java programs:

- You can collect Java trace events using the iSeries Java virtual machine.
- To determine the time that is spent in each Java method, use Java call traces.
- Java profiling locates the relative amount of CPU time that is spent in each Java method and all system functions that are in use by your Java program.
- Use the Java Performance Data Collector to provide profile information about the programs that run on the iSeries server.

Any job session can start and end PEX. Normally, the data that is collected is system wide and pertains to all jobs on the system, including your Java programs. At times, it may be necessary to start and stop the performance collection from inside a Java application. This reduces the collection time and may reduce the large volume of data that is usually produced by a call or return trace. PEX cannot run from within a Java thread. To start and stop a collection, you need to write a native method that communicates to an independent job through a queue or shared memory. Then, the second job starts and stops the collection at the appropriate time.

In addition to application-level performance data, you can use existing iSeries system level performance tools. These tools report statistics on a Java thread basis.

For examples of PEX reports, see the Performance Tools for iSeries, SC41-5340 book. 

Java runtime performance considerations

To greatly improve the runtime performance of your Java^(TM) code, use the Create Java Program (CRTJVAPGM) command before running Java class files, JAR files, or ZIP files. The CRTJVAPGM command uses the bytecodes to create a Java program, which contains optimized native instructions for the iSeries server, and associates the Java program object with the class file, JAR file, or ZIP file.

If the CRTJVAPGM command is not used before running a Java class file, JAR file, or ZIP file, then the Java code runs more slowly the first time, because an optimized Java program at optimization level 10 is created. Subsequent runs are much faster because the Java program is saved and remains associated with the class file or JAR file. Running the bytecodes interpretively may provide acceptable performance during application development, but you may want to use the CRTJVAPGM command before running the Java code in a production environment.

A Just-In-Time compiler improves performance by compiling code, which is optimized to a particular Java virtual machine runtime environment. You should compare the benefits of not having to do CRTJVAPGM with slightly slower program startup and runtime processing that comes with the use of a JIT to decide which optimization technique fits your needs.

If your program is running slowly, enter the Display Java Program (DSPJVAPGM) command to view the attributes of a Java program. >>

Cache class loaders

You can improve the performance of user class loaders by allowing the JVAPGMs created by user class loaders to be cached for reuse. This option is not enabled unless you specify an enabling Java property in one of the following:

- /QIBM/UserData/Java400/SystemDefault.properties
- /home//SystemDefault.properties
- On the RUNJVA CL command
- On the JAVA QSH command

The following property is used to enable this function:

os400.define.class.cache.file

The value of the property should be the name (with the full path) of a valid Java ARchive (JAR) file. The JAR file must contain a valid JAR directory (as built by the jar QSH command), but need to have no other contents other than the single member required to make the jar command function. This JAR file should not be included in any Java classpath.

For example, add the following line to /QIBM/UserData/Java400/SystemDefault.properties:

```
os400.define.class.cache.file=/QIBM/ProdData/Java400/QDefineClassCache.jar
```

The Developer Kit for Java installs a suitable JAR file as /QIBM/ProdData/Java400/QDefineClassCache.jar and installs an example properties file as /QIBM/ProdData/Java400/SystemDefaultCacheExample.properties. You may enable the caching function globally by copying the SystemDefaultCacheExample.properties file to /QIBM/UserData/Java400/ and renaming it to SystemDefault.properties.

Note: Ensure that there is no existing SystemDefault.properties file in that directory before making this change. Alternatively, to affect only a single user's environment, the file may be copied to /home// and similarly renamed.

DSPJVAPGM can be used on the JAR file specified for caching to determine how many JVAPGMs are cached. The **Java programs** field of the DSPJVAPGM display indicates how many JVAPGMs are cached, and the **Java program size** field indicates how much storage is consumed by cached JVAPGMs. Other fields of the display are meaningless when DSPJVAPGM is applied to a JAR file used for caching.

In addition to the os400.define.class.cache.file Java property that enables caching, two other properties can be specified to control the caching characteristics.

os400.define.class.cache.hours

This property can be used to specify how long (in hours) a JVAPGM should persist in the cache.

os400.define.class.cache.maxpgms

This property can be used to specify the maximum number of JVAPGMs that may be cached, with the oldest JVAPGMs being deleted first when this limit is exceeded. The time value used to determine the age of a cached JVAPGM is updated whenever the JVAPGM is referenced.

Examples:


```
os400.define.class.cache.hours=48
os400.define.class.cache.maxpgms=10000
```

The default value for `os400.define.class.cache.hours` is 168 hours (one week), and the maximum value is 9999. The default value for `os400.define.class.cache.maxpgms` is 5000, and the maximum value is 40000. If a value of zero is specified for one of these properties or if the value cannot be parsed as a valid decimal number, then the default value is used.

See Systems properties for more information on the properties previously mentioned for caching. [«](#)

Select which mode to use when running a Java program

When you run a Java^(TM) program, you can select which mode you would like to use. All modes verify the code and create a Java program object to hold the preverified form of the program. You can use any of the following modes:

- Interpreted
- Direct processing
- Just-In-Time (JIT) compile
- Just-In-Time (JIT) compile and direct processing

Selection mode	Details
Interpreted	Each bytecode is interpreted at runtime. For information on running your Java program in the interpreted mode, see the Run Java (RUNJVA) command.
Direct Processing	Machine instructions for a method are generated during the first call to that method, and saved for use the next time that the program runs. One copy is also shared for the entire system. For information on running your Java program using direct processing, see the Run Java (RUNJVA) command.

Selection mode	Details
Just-In-Time (JIT) compile	<p>Machine instructions for a method are generated during the first call to that method, and saved for the duration of the Java virtual machine run.</p> <p>To use the Just-In-Time compiler, you need to set the compiler value to <code>jitc</code>. You can set the value by adding an environment variable or setting the <code>java.compiler</code> system property. Select one method from the list below to set the compiler value:</p> <ul style="list-style-type: none"> From a command line prompt on your iSeries server, add the environment variable by using the Add Environment Variable (ADDENVVAR) command. Then, run your Java program using the Run Java (RUNJVA) command or JAVA command. For example, use: ADDENVVAR ENVVAR (JAVA_COMPILER) VALUE(jitc) JAVA CLASS(Test) Set the <code>java.compiler</code> system property on the iSeries command line. For example, enter <code>JAVA CLASS(Test) PROP((java.compiler jitc))</code> Set the <code>java.compiler</code> system property on the Qshell Interpreter command line. For example, enter <code>java -Djava.compiler=jitc Test</code> <p>Once you set this value, the JIT compiler optimizes all of the Java code before running it.</p>
Just-In-Time (JIT) compile and Direct Processing	<p>The most common way to use the Just-In-Time (JIT) compiler is with the <code>jitc_de</code> option. When running with this option, programs that have already been optimized with direct processing run in direct processing mode. Programs that have not been optimized for direct optimization run in JIT mode.</p> <p>To use JIT and direct processing together, you need to set the compiler value to <code>jitc_de</code>. You can set the value by adding an environment variable or setting the <code>java.compiler</code> system property. Select one method from the following list to set the compiler value:</p> <ul style="list-style-type: none"> Add the environment variable by entering the Add Environment Variable (ADDENVVAR) command on the iSeries command line. Then, run your Java program using the Run Java (RUNJVA) command or JAVA command. For example, enter ADDENVVAR ENVVAR (JAVA_COMPILER) VALUE(jitc_de) JAVA CLASS(Test) Set the <code>java.compiler</code> system property on the iSeries command line. For example, enter <code>JAVA CLASS(Test) PROP((java.compiler jitc_de))</code> Set the <code>java.compiler</code> system property on the Qshell Interpreter command line. For example, enter <code>java -Djava.compiler=jitc_de Test</code> <p>Once this value is set, the Java program for the class file that was created as direct processing is used. If the Java program was not created as direct processing, the class file is optimized by the JIT prior to running. For more information, see Comparison of Just-In-Time compiler and direct processing</p>

There are three ways that you can run a Java program (CL, QSH, and JNI). Each has a unique way to specify the mode. This table shows how that is done.

Mode	CL Command	QShell Command	JNI Invocation API
Interpret	INTERPRET(*YES)	-Djava.compiler=NONE -interpret	os400.run.mode="interpret"
DE	INTERPRET(*NO)	-Djava.compiler=NONE	<ul style="list-style-type: none"> os400.run.mode="program_created=pc" os400.create.type="direct"
JIT	INTERPRET(*JIT)	-Djava.compiler="jitc"	os400.run.mode="jitc"
JIT_DE(default)	INTERPRET(*OPTIMIZE) OPTIMIZE(*JIT)	-Djava.compiler="jitc_de"	os400.run.mode="jitc_de"

Java interpreter

The Java^(TM) interpreter is the part of the Java virtual machine that interprets Java class files for a particular hardware platform. The Java interpreter decodes each bytecode and runs a series of machine instructions for that bytecode.

Static compilation

The Java^(TM) transformer is an IBM Operating System/400 (OS/400) component that preprocesses class files to prepare them to run using the iSeries Java virtual machine. The Java transformer creates an optimized program object that is persistent and is associated with the class file. In the default case, the program object contains a compiled, 64-bit RISC machine instruction version of the class. The Java interpreter does not interpret the optimized program object at runtime. Instead, it directly runs when the class file is loaded.

Java programs are optimized using the JIT by default. To use the Java transformer, you either do CRTJVAPGM, or specify the use of the transformer on the RUNJVA or JAVA command.

You can use the Create Java Program (CRTJVAPGM) command to explicitly start the Java transformer. The CRTJVAPGM command optimizes the class file or JAR file while the command runs, so nothing needs to be done while the program is running. This improves the speed of the program the first time that it runs. Using the CRTJVAPGM command, instead of relying on default optimization, ensures the best optimization possible and also improves the use of space for the Java programs that are associated with the class file or JAR file.

Using the CRTJVAPGM command on a class file, JAR file, or ZIP file causes all the classes in the file to be optimized, and the resulting Java program object are persistent. This results in better runtime performance. You can also change the optimization level or select an optimization level other than the default of 10 by using the CRTJVAPGM command or the Change Java Program (CHGJVAPGM) command. At optimization level 40, interclass binding is performed between the classes within a JAR file, and in some cases, the classes are inlined. Interclass binding improves the call speed. Inlining removes the overhead of a method call entirely. In some cases, you can inline methods between classes within the JAR file or ZIP file. Specifying OPTIMIZE(*INTERPRET) on the CRTJVAPGM command causes any classes that are specified on the command to be verified and prepared to run in interpreted mode.

The Run Java (RUNJVA) command can also specify OPTIMIZE(*INTERPRET). This parameter specifies that any classes running under the Java virtual machine are interpreted, regardless of the optimization level of the associated program object. This is useful when debugging a class that was transformed with an optimization level of 40. To force interpretation, use INTERPRET(*YES).

➤ See Cache class loaders for information on reusing your Java programs created by class loaders. ⏪

Java static compilation performance considerations

You can determine the speed of transformation by the optimization level that you set. Optimization level 10 transforms the fastest, but the resulting program is generally slower than one set at a higher optimization level. Optimization level 40 takes longer to transform, but is likely to run faster.

A small number of Java^(TM) programs may not optimize to level 40. Thus, a few programs that do not run at level 40, may run at level 30 instead. You can run programs that do not run at optimization level 40 by using licensed internal code optimization LICOPT parameter strings. However, performance at level 30 may be sufficient for your program.

If you are having problems running Java code that seemed to work on another Java virtual machine, try using optimization level 30 instead of level 40. If this works, and your performance is acceptable, you do not need to do anything else. If you need better performance, see LICOPT parameter strings for information on how to enable and disable various forms of optimization. For example, you could first try creating the program using OPTIMIZE(40) LICOPT(NoPreresolveExtRef). If your application contains dead calls to classes that are not available, this LICOPT value allows your program to run without problems.

To determine what level of optimization your Java programs were created at, you can use the Display Java Program (DSPJVAPGM) command. To change the optimization level of your Java program, use the Create Java Program (CRTJVAPGM) command.

Just-In-Time compiler

A Just-In-Time (JIT) compiler is a platform-specific compiler that generates machine instructions for each method upon the first call to that method. To improve performance over interpretation, the JIT compiler compiles code as you need it.

To understand the difference between the JIT compiler and direct processing, see Comparison of the JIT compiler and direct processing.

Refer to Java^(TM) system properties for information on the java.compiler property. Select the version you are running.

Comparison of Just-In-Time compiler and direct processing

If you are trying to decide whether to use the Just-In-Time compiler or direct processing mode to run your Java^(TM) program, this table provides additional information to help you make the best choice for your situation.

Just-In-Time compiler	Direct processing
Provides an automatic compilation of any method when needed. This method is much faster than direct processing.	You compile an entire class or JAR file using the Create Java Program (CRTJVAPGM) command or the files are compiled automatically at runtime.
Used during program development to cut optimization costs while the code is changing. Also used in deployment of small or relatively low usage applications. It can also be used with highly dynamic applications where code is generated or discovered at runtime.	Used for relatively large applications. Direct processing code, at optimization level 40, is typically faster than JIT. Most ready-to-deploy server applications use direct processing at optimization level 40, because they are likely to be in use by multiple users at any given time, and the cost of using JIT over and over is too much overhead.
Optimizations are limited to those that can be performed rapidly at runtime.	More complex optimizations are possible, since optimization is not performed at runtime.

Optimization levels

By entering a value in the **Optimization level** field, you specify the optimization level of the Java^(TM) program that attaches to the class file object or JAR file object. You can control the size and performance

of the Java program by using this option. For the optimization levels listed, the internal form includes iSeries machine instructions. These machine instructions are optimized based on the specified optimization level. The iSeries server directly runs the machine instructions when the program runs.

This list shows how the optimization levels differ and what they do:

10

The Java program contains a transformed version of the class file bytecodes, but has only minimal additional compiler optimization. You can display and change variables while debugging. For V5R1 and higher, an optimization level of 10 may unnecessarily lengthen the time it takes to compile your program. Instead, use an optimization level of at least 20.

20

The Java program contains a compiled version of the class file bytecodes and performs additional compiler optimization. You can display, but not change variables while debugging.

30

The Java program contains a compiled version of the class file bytecodes and performs more compiler optimization than optimization level 20. Debugging at higher optimization levels is more difficult, because optimization reduces the ability to stop at precise points in the code and display program variables. You can display, but not change variables while debugging. The values that are presented may not be the current value of the variable.

40

The Java program contains a compiled version of the class file bytecodes and performs more compiler optimization than optimization level 30. Optimization level 40 includes cross-class optimizations. In a small number of cases, the order in which static initializers are run for unrelated classes (not related by inheritance nor containment) may be different than outlined in the static initialization specification. In addition, it includes optimization that disables call and instruction tracing.

Note: If your Java program fails to optimize or throws an exception at optimization level 40, use optimization level 30.

Java garbage collection

Garbage collection is the process of freeing storage that is used by objects that are no longer referred to by a program. With garbage collection, programmers no longer have to write error prone code to explicitly "free" or "delete" their objects. This code frequently results in "memory leak" program errors. The garbage collector automatically detects an object or group of objects that the user program can no longer reach. It does this because there are no references to that object in any program structure. Once the object has been collected, you can allocate the space for other uses.

The Java^(TM) runtime environment includes a garbage collector that frees memory that is no longer in use. The garbage collector runs automatically, as needed.

The garbage collector can also be started explicitly under the control of the Java program using the `java.lang.Runtime.gc()` method.

See IBM Developer Kit for Java advanced garbage collection for IBM Developer Kit for Java specifics.

IBM Developer Kit for Java advanced garbage collection

The IBM Developer Kit for Java^(TM) implements an advanced garbage collector algorithm. This algorithm allows the discovery and collection of unreachable objects without significant pauses in the operation of the Java program. A concurrent collector cooperatively discovers the references to objects under the running threads, instead of a single thread.

Many garbage collectors are "stop-the-world". This means that at the point where a collection cycle occurs, all threads, except the thread that does garbage collection, stop while the garbage collector does its work. When this happens, Java programs experience a pause, and any multiple processor capability of the platform is wasted relative to Java, while the collector does its work. The iSeries algorithm does not stop all program threads simultaneously. It allows those threads to continue operation while the garbage collector completes its task. This prevents the pauses, and allows all processors to be used during garbage collection.

Garbage collection occurs automatically based on parameters that you specify when you start the Java virtual machine. Garbage collection can also be started explicitly under the control of the Java program by using the `java.lang.Runtime.gc()` method.

For a basic definition, see Java garbage collection.

Java garbage collection performance considerations

Garbage collection on the iSeries Java^(TM) virtual machine operates in a continuous asynchronous mode. The garbage collection-initial size (GCHINL) parameter on the Run Java (RUNJVA) command may affect application performance. This parameter specifies the amount of new object space that is allowed between garbage collections. A small value may cause too much garbage collection overhead. A large value may limit garbage collection and cause out of memory errors. However, for most applications, the default values should be correct.

Garbage collection determines that an object is no longer needed by evaluating whether there are any valid references to that object.

Java Native Method Invocation performance considerations

Native method invocation on an iSeries server may not perform as well as native method invocation on other platforms. Java^(TM) on the iSeries server has been optimized by moving the Java virtual machine below the machine interface (MI). Native method invocation requires a call to above MI code and may require expensive Java Native Interface (JNI) calls back into the Java virtual machine. Native methods should not carry out small routines, which you can easily write in Java. Only use native methods to start system functions that are relatively long running and are not available directly in Java.

Java method inlining performance considerations

Method inlining can significantly improve method call performance. Any method that is final is a potential candidate for inlining. The inline feature is available on the iSeries server through the `javac -o` option at compilation time. The size of your class files and transformed Java^(TM) program increases if you use the `javac -o` option. You should consider both the space and performance characteristics of your application when using the `-o` option.

The Java transformer enables inlining for optimization level 30 and optimization level 40. Optimization level 30 enables some inlining of final methods within a single class. Optimization level 40 enables inlining of final methods within a ZIP file or JAR file. You can control method inlining with the `AllowInlining` and `NoAllowInlining` LICOPT parameter strings. The iSeries interpreter does not perform method inlining.

Java exception performance considerations

The iSeries exception architecture allows versatile interrupt and retry capabilities. It allows mixed language interaction. Throwing Java^(TM) exceptions on an iSeries server may be more expensive than on other platforms. This should not affect overall application performance unless Java exceptions are routinely used in the normal application path.

Java call trace performance tools

Java^(TM) method call traces provide significant performance information about the time that is spent in each Java method. On other Java virtual machines, you may have used the `-prof` (profiling) option on the `java` command. To enable method call tracing on an iSeries server, you must specify the Enable Performance Collection (ENBPFCOL) command on the Create Java Program (CRTJVAPGM) command line. After creating your Java program with this keyword, you can start the collection of method call traces by using a Performance Explorer (PEX) definition that includes the call/return trace type.

Call/return trace output produced with the Print Performance Explorer Report (PRTPEXRPT) command shows the central processing unit (CPU) time for each call for every Java method that is traced. In some cases, you may not be able to enable all of the class files for call return tracing. Or, you may be calling native methods and system functions that are not enabled for tracing. In this situation, all of the CPU time that is spent in these methods or system functions accumulates. Then, it is reported to the last Java method that is called and has been enabled.

Java event trace performance tools

The iSeries Java^(TM) virtual machine enables the trace of certain Java events. These events can be collected without any instrumentation in the Java code. These events include activities, such as garbage collection, thread creation, class loading, and locking. The Run Java (RUNJVA) command does not specify these events. Instead, you create a Performance Explorer (PEX) definition and use the Start Performance Explorer (STRPEX) command to collect the events. Each event contains useful performance information, such as time stamp and central processing unit (CPU) cycles. You can trace both Java events and other system activities, such as disk input and output, with the same trace definition.

For a complete description of the Java events, see the Performance Tools for iSeries, SC41-5340 book.



Java profiling performance tools

System wide central processing unit (CPU) profiling calculates the relative amount of CPU time that is spent in each Java^(TM) method and all system functions in use by your Java program. Use a Performance Explorer (PEX) definition that traces performance monitor counter overflow (*PMCO) run cycle events. Samples are typically specified in intervals of one millisecond. To collect a valid trace profile, you should run your Java application until it accumulates two to three minutes of CPU time. This should produce over 100,000 samples. The Print Performance Explorer Report (PRTPEXRPT) command produces a histogram of CPU time that is spent across the entire application. This includes every Java method and all system-level activity. The Performance Data Collector (PDC) tool also provides profile information about the programs that run on the iSeries server.

Note: CPU profiling does not show relative CPU usage for Java programs that are interpreted.

Java Virtual Machine Profiler Interface

➤ The Java^(TM) Virtual Machine Profiler Interface (JVMPi) is an experimental interface for profiling the Java virtual machine (JVM), which was first disclosed and implemented in Sun's Java 2 SDK, Standard Edition (J2SDK), version 1.2.

JVMPI support places hooks in the JVM and the Just-in-time (JIT) compiler, which when activated, provide event information to a profiling agent. The profiling agent is implemented as a dynamic link library (DLL). The profiler sends control information to the JVM for enabling and disabling JVMPI events. For example, the profiler may not be interested in method Entry or Exit hooks and could tell the JVM that it does not want to receive these event notifications. The JVM and JIT have JVMPI event hooks embedded that send event notifications to the profiling agent if the event is enabled. The profiler tells the JVM which events are of interest and the JVM sends notifications of the events to the profiler when they occur.

For more information, see JVMPI  by Sun Microsystems, Inc. 

Collect Java performance data

To collect Java^(TM) performance data on an iSeries server, follow these steps:

1. Create a Performance Explorer (PEX) definition that specifies:
 - A user-defined name
 - Type of data collection
 - Job name
 - Series of system events that you would like to collect system information about

Note: A PEX definition of *STATS is preferable to a *TRACE definition if the output that you want is the `java_g -prof` type, and you know the specific job name of the Java program.

Here is an example of a *STATS definition:

```
ADDPEXDFN DFN(YOURDFN) JOB(*ALL/YOURID/QJVACMSRV) DTAORG(*HIER)
TEXT('your stats definition')
```

This *STATS definition does not get all Java events running. Only the Java events that are in your own Java session are profiled. This mode of operation may increase the time that it takes to run the Java program.

Here is an example of a *TRACE definition:

```
ADDPEXDFN DFN(YOURDFN) TYPE(*TRACE) JOB(*ALL) TRCTYPE(*SLTEVT)
SLTEVT(*YES) PGMEVT(*JVAENTRY *JVAEXIT)
```

This *TRACE definition collects any Java entry event and exit event from any Java program in the system that you create with `ENBPFRCOL(*ENTRYEXIT)`. This causes the analysis of this type of collection to be slower than a *STATS trace, depending on how many Java program events you have and the duration of the PEX data collection.

2. Enable the *JVAENTRY and *JVAEXIT, under the program events category on the PEX definition, so that PEX recognizes the Java entry and exits.

Note: If you are running the Java code using the Just-in-time (JIT) compiler, you do not enable entry and exit as you would if you were using the `CRTJVAPGM` command for direct processing. Instead, JIT generates code with entry and exit hooks when you use the `os400.enbprfcol` system property.

3. Prepare the Java program to report program events to the iSeries Performance Data Collector. You can do this by using the Create Java Program (`CRTJVAPGM`) command on any Java program that you want to report performance data on. You must create the Java program by using the `ENBPFRCOL(*ENTRYEXIT)` parameter.

Note: You must repeat this step for every Java program that you want to collect performance data on. If you do not perform this step, no performance data is collected by the PEX and no output is produced by running the Java Performance Data Converter (JPDC) tool.

4. Start the PEX data collection by using the Start Performance Explorer (`STRPEX`) command.
5. Run the program that you would like to analyze. This program should not be in a production environment. It generates a large amount of data in a small amount of time. You should limit the

collection time to five minutes. A Java program that runs for this amount of time generates a lot of PEX system data. If too much data is collected, an unreasonable amount of time is required to process it.

6. End the PEX data collection by using the End Performance Explorer (ENDPEX) command.

Note: If this is not the first time that you have ended PEX data collection, you must specify a replace file of *YES or it does not save your data.

7. Run the JPDC tool.
8. Connect the integrated file system directory to the system with the viewer of your choice: java_g -prof viewer or Jinsight viewer. You can copy this file from you iSeries server and use it as input to any suitable profiling tool.

Performance Data Collector tool

The Performance Data Collector (PDC) tool provides profile information about the programs that run on the iSeries server.

The industry-standard profile option on many Java^(TM) virtual machines depends on the implementation of the java_g feature. This is a special debug version of the Java virtual machine, which offers the -prof option. You specify this option on a call to a Java program. When you specify this option, the Java virtual machine produces a record file that contains information about which parts of the Java program were operating during the duration of the program. The Java virtual machine generates this information in real time.

On the iSeries server, the Performance Explorer (PEX) feature analyzes programs and record-specific system events. A DB2^(R) database stores this information and retrieves it using SQL functions. PEX information is the repository for specific program information that produces Java profile data. This profile data is compatible with java_g -prof program profile information. The Java Performance Data Converter (JPDC) tool provides java_g -prof program output and program profile information for a specific IBM tool, which is known as Jinsight.

For information on how collect Java performance data, see Collect Java performance data.

Java Performance Data Converter tool

The Java^(TM) Performance Data Converter (JDPC) tool provides a way for you to create Java performance data about the Java programs that are running on your iSeries server. This performance data is compatible with the performance data output of Sun Microsystems, Inc.'s Java virtual machine java_g -prof option and IBM Jinsight output.

Note: The JDPC tool does not produce readable output. Use a Java profiling tool that accepts java_g -prof or Jinsight data to analyze your data.

The JDPC tool accesses the iSeries Performance Explorer (PEX) data that DB2/400 (using JDBC) stores. It converts the data to either Jinsight or general performance types. Then, JDPC stores the output file in the integrated file system at a user-specified location.

Note: You must follow appropriate iSeries PEX data collection procedures to collect PEX data while running your specified Java application on an iSeries server. You must set a PEX definition with defines the entrance and exit of a program or a collect and store procedure. For details on how collect PEX data

and set a PEX definition, see the Performance Tools for iSeries, SC41-5340 book. 

For information on how to run JPDC, see Run the Java Performance Data Converter.

You can start the JPDC program by using either the Qshell command line interface or Run Java (RUNJVA) command.

Run the Java Performance Data Converter

To run the Java^(TM) Performance Data Converter (JPDC) for performance data collection, follow these steps.

1. Enter the first input argument, which is either `general` for the `java_g -prof` or `jinsight` for Jinsight output.
2. Enter the second input argument, which is the name of the Performance Explorer (PEX) definition that was used to collect the data.
Note: You should restrict this name to four or five characters, because of the internal use of connections of this name.
3. Enter the third input argument, which is the name of the file that the JPDC tool generates. This generated file writes to your current integrated file system directory. Use the `cd (PF4)` command to specify an integrated file system current directory.
4. Enter the fourth input argument, which is the name of the iSeries host relational database directory entry. Use the `Work with Relational Database Directory Entry (WRKRDBDIRE)` command to see what the name is. It is the only relational database where the `*LOCAL` is indicated.

To operate this code the `/QIBM/ProdData/Java400/ext/JPDC.jar` file must be in the Java classpath on the iSeries server. When the program is done running, a text output file can be found in the current directory.

You can run JPDC by using the iSeries command line or Qshell environment. See `Example: Run the Java Performance Data Converter` for details.

Example: Run the Java Performance Data Converter

You can either use the iSeries command line or the Qshell environment to run the Java^(TM) Performance Data Converter (JPDC).

Using the iSeries command line:

1. Enter the `Run Java (RUNJVA)` command or `JAVA` command on the iSeries command line.
2. Enter `com.ibm.as400.jpdc.JPDC` on the class parameter line.
3. Enter `general pexdfn mydir/myfile myrdbdire` on the parameter line.
4. Enter `'/QIBM/ProdData/Java400/ext/JPDC.jar'` on the classpath parameter line.

Note: You can omit the classpath if the `'/QIBM/ProdData/Java400/ext/JPDC.jar'` string is in the `CLASSPATH` environment variable. You can use either the `Add Environment Variable (ADDENVVAR)` command, `Change Environment Variable (CHGENVVAR)` command, or `Work with Environment Variable (WRKENVVAR)` command to add this string to the `CLASSPATH` environment variable.

Using the Qshell environment:

1. Enter the `Start Qshell (STRQSH)` command to start the Qshell Interpreter.
2. Enter this on the command line:

```
java -classpath /QIBM/ProdData/Java400/ext/JPDC.jar com.ibm.as400/jpdc/JPDC
jinsight pexdfn mydir/myfile myrdbdire
```

Note: You can omit the classpath if the `'/QIBM/ProdData/Java400/ext/JPDC.jar'` string is added to your current environment. You can use either the `ADDENVVAR` command, `CHGENVVAR`, or `WRKENVVAR` command to add this string to your current environment.

For background information, see `Run the Java Performance Data Converter`.

Chapter 6. Commands and tools for the IBM Developer Kit for Java

When using the IBM Developer Kit for Java^(TM), you can either use Java tools with the Qshell Interpreter or CL commands.

If you have prior Java programming experience, you may be more comfortable using the Qshell Interpreter Java tools, because they are similar to the tools that you would use with Sun Microsystems, Inc. Java Development Kit. See Qshell Interpreter for information about using the Qshell environment.

If you are an iSeries programmer, you may want to use the CL commands for Java that are typical to the iSeries server environment. Read on for more information about using CL commands and iSeries Navigator commands.

You can use any of these commands and tools with the IBM Developer Kit for Java:

- The Qshell environment includes the Java development tools that are typically required for program development.
- The CL environment contains the CL commands for optimizing and managing Java programs.
- The iSeries Navigator commands also create and run optimized Java programs.

Java tools that are supported by the IBM Developer Kit for Java

The IBM Developer Kit for Java^(TM) supports these tools:

- Java ajar tool
- Java appletviewer tool
- Java extcheck tool
-  Java idlj tool 
- Java jar tool
- Java jarsigner tool
- Java javac tool
- Java javadoc tool
- Java javah tool
- Java javakey tool
- Java javap tool
- Java keytool
- Java native2ascii tool
- Java policytool
- Java rmic tool
- Java rmid tool
- Java rmiregistry tool
- Java serialver tool
- Java tnameserv tool

With a few exceptions, the Java tools, except the ajar tool, support the syntax and options that are documented by Sun Microsystems, Inc. They must all run by using the Qshell Interpreter.

You can start the Qshell Interpreter by using the Start Qshell (STRQSH or QSH) command. When the Qshell Interpreter is running, a QSH Command Entry display appears. All output and messages from Java

tools and programs that run under Qshell appear in this display. Any input to a Java program is also read from this display. See Java command in Qshell for more details.

Note: Functions of iSeries command entry are not available directly from within the Qshell. To get a command entry, press F21 (CL command entry).

Java tools

- Java ajar tool
- Java appletviewer tool
 - Run the Java appletviewer tool with Remote Abstract Window Toolkit
- Java extcheck tool
- Java idlj tool
- Java jar tool
- Java jarsigner tool
- Java javac tool
- Java javadoc tool

Java ajar tool

The ajar tool is an alternative interface to the jar tool that you use to create and manipulate Java^(TM) ARchive (JAR) files. You can use the ajar tool to manipulate both JAR files and ZIP files.

If you need a ZIP interface or UNZIP interface, use the ajar tool instead of the jar tool.

The ajar tool lists the contents of JAR files, extracts from JAR files, creates new JAR files, and supports many of the ZIP formats just as the jar tool does. Additionally, the ajar tool supports adding and deleting files in existing JAR files.

The ajar tool is available using the Qshell Interpreter. For more details, see ajar - Alternative Java archive.

Java appletviewer tool

The appletviewer tool allows you to run applets without a web browser. It is compatible with the appletviewer tool that is supplied by Sun Microsystems, Inc.

The appletviewer tool is available using the Qshell Interpreter. You need to use Remote Abstract Window Toolkit to run the appletviewer tool. See Run the Java appletviewer tool with Remote AWT for information about how to set up Remote AWT to use the appletviewer tool.

For more information about the appletviewer tool, see the appletviewer tool by Sun Microsystems, Inc.

Run the Java appletviewer tool with Remote Abstract Window Toolkit: To use the Java appletviewer tool, you must set up the Remote Abstract Window Toolkit for Java on a Windows^(R) remote display, and use either the sun.applet.AppletViewer class or run the appletviewer tool in the Qshell Interpreter with Remote AWT properties.

For example, if you use the sun.applet.AppletViewer class and run example1.html out of the TicTacToe directory, the command line should look like this:

```
JAVA CLASS(sun.applet.AppletViewer) PARM('example1.html') CLASSPATH('/TicTacToe')
PROP((RmtAwtServer '1.1.11.11') (os400.class.path.rawt 1)(java version 1.3))
```

If you use the appletviewer tool in the Qshell Interpreter and run example1.html out of the TicTacToe directory, the commands should look like this:

```
qsh "enter"
cd TicTacToe "enter"
Appletviewer -J-DRmtAwtServer=1.1.11.11 -J-Dos400.class.path.rawt=1 -J-Djava.version=1.3 example1.html
```

Note: -J are runtime flags for Appletviewer. -D are properties.

Java extcheck tool

In Java 2 SDK (J2SDK), Standard Edition, version 1.2 and higher, the extcheck tool detects version conflicts between a target JAR file and currently installed extension JAR files. It is compatible with the keytool that is supplied by Sun Microsystems, Inc.

The extcheck tool is available using the Qshell Interpreter.

For more information about the extcheck tool, see the extcheck tool by Sun Microsystems, Inc.

Java idlj tool

➤ The idlj tool generates Java bindings from a given Interface Definition Language (IDL) file. The idlj tool is also known as the IDL-to-Java compiler. It is compatible with the idlj tool that is supplied by Sun Microsystems, Inc. This tool only works for Java Development Kits 1.3 and 1.4.

For more information about the idlj tool, see the idlj tool by Sun Microsystems, Inc. ⏪

Java jar tool

The jar tool combines multiple files into a single Java ARchive (JAR) file. It is compatible with the jar tool that is supplied by Sun Microsystems, Inc.

The jar tool is available using the Qshell Interpreter.

For an alternative interface to the jar tool, see the ajar tool for creating and manipulating JAR files.

For more information about iSeries file systems, see the Integrated file system or Files in the integrated file system.

For more information about the jar tool, see the jar tool by Sun Microsystems, Inc.

Java jarsigner tool

In Java 2 SDK (J2SDK), Standard Edition, version 1.2 and higher, the jarsigner tool signs JAR files and verifies signatures on signed JAR files. The jarsigner tool accesses the keystore, which the keytool creates and manages, when it needs to find the private key for signing a JAR file. In J2SDK, the jarsigner and keytool tools replace the javakey tool. It is compatible with the jarsigner tool that is supplied by Sun Microsystems, Inc.

The jarsigner tool is available using the Qshell Interpreter.

For more information about the jarsigner tool, see the jarsigner tool by Sun Microsystems, Inc.

Java javac tool

The javac tool compiles Java programs. It is compatible with the javac tool that is supplied by Sun Microsystems, Inc. with one exception.

-classpath

Does not override the default classpath. Instead, it is appended to the system default classpath. The -classpath option does override the CLASSPATH environment variable.

The javac tool is available using the Qshell Interpreter.

If you have JDK 1.1.x installed on your iSeries server as your default, but you need to run the java command from version 1.2 or higher, enter this command:

```
javac -djava.version=1.2 <my_dir> MyProgram.java
```

For more information about the javac tool, see the javac tool by Sun Microsystems, Inc.

Java javadoc tool

The javadoc tool generates API documentation. It is compatible with the javadoc tool that is supplied by Sun Microsystems, Inc.

The javadoc tool is available using the Qshell Interpreter.

For more information about the javadoc tool, see the javadoc tool by Sun Microsystems, Inc.

Java tools

- Java javah tool
- Java javakey tool
- Java javap tool
- Java keytool
- Java native2ascii tool
- Java policytool
- Java rmic tool
- Java rmid tool
- Java rmiregistry tool
- Java serialver tool
- Java tnameserv tool

Java javah tool

The javah tool facilitates the implementation of Java^(TM) native methods. It is compatible with the javah tool that is supplied by Sun Microsystems, Inc. with a few exceptions.

Note: Writing native methods means that your application is not 100% pure Java. It also means that your application is not directly portable across platforms. Native methods are, by nature, platform or system-specific. Using native methods may increase your development and maintenance costs for your applications.

The javah tool is available using the Qshell Interpreter. It reads a Java class file and creates a C-language header file in the current working directory. The header file that is written is an iSeries Stream File (STMF). It must be copied to a file member before it can be included in a C program on the iSeries server.

The javah tool is compatible with the tool that is provided by Sun Microsystems, Inc. If these options are specified, however, the iSeries server ignores them.

-td The javah tool on the iSeries server does not require a temporary directory.

-stubs Java on the iSeries server only supports the Java Native Interface (JNI) form of native methods. Stubs were only required for the pre-JNI form of native methods.

-trace Relates to the .c stub file output, which Java on the iSeries server does not support.

-v Not supported.

Note: The -jni option must always be specified. The iSeries server does not support native method implementations prior to JNI.

For more information about the javah tool, see the javah tool by Sun Microsystems, Inc.

Java javakey tool

Use the javakey tool for encryption key, and certificate generation and management, including generation of digital signatures for applets. It is compatible with the javakey tool that is supplied by Sun Microsystems, Inc. in Java Development Kit (JDK) versions 1.1.x.

In Java 2 Software Development Kit (J2SDK), Standard Edition, version 1.2 or higher, the javakey tool is obsolete. Due to bugs in JDK versions 1.1.x, code that is signed using the 1.1.x javakey tool is recognized as unsigned in the J2SDK, version 1.2 or higher. If you sign code using J2SDK, version 1.2 or higher, it is recognized as unsigned on JDK 1.1.x versions.


Note: The iSeries secure sockets layer (SSL) support cannot access keys created by this tool. Instead, you must use the certificate and key containers that are integrated into the iSeries server and created or imported with the Digital Certificate Manager. See *Make a Java application secure with the secure sockets layer* for more information.

Applet packaging and applet signing is dependent on your browser. Check your browser documentation to ensure that your browser is compatible with the Java JAR file format and javakey applet signing.

Note: The files that are created by the javakey tool contain sensitive information. Appropriate integrated file system security measures protect the public and private key files.

The javakey tool is available using the Qshell Interpreter.

For more information about iSeries file systems, see the *Integrated file system* or *Files in the integrated file system*.

For more information about the javakey tool, see the *javakey tool by Sun Microsystems, Inc.* 

Java javap tool

The javap tool disassembles compiled Java files and prints out a representation of the Java program. This may be helpful when the original source code is no longer available on a system.

It is compatible with the javap tool that is supplied by Sun Microsystems, Inc. with a few exceptions.

- b** This option is ignored. Backward compatibility is not required, because Java on the iSeries server only supports Java Development Kit (JDK) 1.1.4 and later.
- p** On the iSeries server, -p is not a valid option. You must spell out -private.
- verify** This option is ignored. The javap tool does not do verification on the iSeries server.

The javap tool is available using the Qshell Interpreter.

Note: The use of the javap tool to disassemble classes may violate the license agreement for those classes. Consult the license agreement for the classes before using the javap tool.

For more information about the javap tool, see the *javap tool by Sun Microsystems, Inc.*

Java keytool

In Java 2 SDK (J2SDK), Standard Edition, version 1.2 or higher, the keytool creates public and private key pairs, self-signed certificates, and manage keystores. In J2SDK, the jarsigner and keytool tools replace the javakey tool. It is compatible with the keytool that is supplied by Sun Microsystems, Inc.

The keytool is available using the Qshell Interpreter.

For more information about the keytool, see the *keytool by Sun Microsystems, Inc.*

Java native2ascii tool

The `native2ascii` tool converts a file with native-encoded characters (characters which are non-Latin 1 and non-Unicode) to one with Unicode-encoded characters. It is compatible with the `native2ascii` tool that is supplied by Sun Microsystems, Inc.

The `native2ascii` tool is available using the Qshell Interpreter.

For more information about the `native2ascii` tool, see the `native2ascii` tool by Sun Microsystems, Inc.

Java policytool

In Java 2 SDK, Standard Edition, the `policytool` creates and changes the external policy configuration files that define the Java security policy of your installation. It is compatible with the `policytool` that is supplied by Sun Microsystems, Inc.

The `policytool` is a graphical user interface (GUI) tool available using the Qshell Interpreter and the Remote Abstract Window Toolkit. See IBM Developer Kit for Java Remote Abstract Window Toolkit for more information.

For more information about the `policytool`, see the `policytool` by Sun Microsystems, Inc.

Java rmic tool

The `rmic` tool generates stub files and class files for Java objects. It is compatible with the `rmic` tool that is supplied by Sun Microsystems, Inc.

The `rmic` tool is available using the Qshell Interpreter.

For more information about the `rmic` tool, see the `rmic` tool by Sun Microsystems, Inc.

Java rmid tool

In Java 2 SDK (J2SDK), Standard Edition, the `rmid` tool starts the activation system daemon, so objects can be registered and activated in a Java virtual machine. It is compatible with the `rmid` tool that is supplied by Sun Microsystems, Inc.

The `rmid` tool is available using the Qshell Interpreter.

For more information about the `rmid` tool, see the `rmid` tool by Sun Microsystems, Inc.

Java rmiregistry tool

The `rmiregistry` tool starts a remote object registry on a specified port. It is compatible with the `rmiregistry` tool that is supplied by Sun Microsystems, Inc.

The `rmiregistry` tool is available using the Qshell Interpreter.

For more information about the `rmiregistry` tool, see the `rmiregistry` tool by Sun Microsystems, Inc.



Java serialver tool

The `serialver` tool returns the version number or serialization-unique identifier for one or more classes. It is compatible with the `serialver` tool that is supplied by Sun Microsystems, Inc.

The `serialver` tool is available using the Qshell Interpreter.

For more information about the `serialver` tool, see the `serialver` tool by Sun Microsystems, Inc.

Java tnameserv tool

In Java 2 SDK (J2SDK), Standard Edition, version 1.3  or higher , the `tnameserv` (Transient Naming Service) tool provides access to the naming service. It is compatible with the `tnameserv` tool that is supplied by Sun Microsystems, Inc.

The tnameserv tool is available using the Qshell Interpreter.

Java command in Qshell

The java command in Qshell runs Java^(TM) programs. It is compatible with the java tool that is supplied by Sun Microsystems, Inc. with a few exceptions.

The IBM Developer Kit for Java ignores these options of the java command in Qshell.

Option	Description
-cs	This option is not supported.
-checksource	This option is not supported.
-debug	This option is supported by the iSeries internal debugger.
-noasyncgc	Garbage collection is always running with the IBM Developer Kit for Java.
-noclassgc	Garbage collection is always running with the IBM Developer Kit for Java.
-prof	The iSeries server has its own performance tools.
-ss	This option is not applicable on the iSeries server.
-oss	This option is not applicable on the iSeries server.
-t	The iSeries server uses its own trace function.
-verify	Always verify on the iSeries server.
-verifyremote	Always verify on the iSeries server.
-noverify	Always verify on the iSeries server.

On the iSeries server, the -classpath option does not override the default classpath. Instead, it is appended to the system default classpath. The -classpath option does override the CLASSPATH environment variable.

The java command in Qshell supports new options for the iSeries server. These are the new options that are supported.

Option	Description
» -chkpath «	This option checks for public write access to directories in the CLASSPATH.
-opt	This option specifies the optimization level.
» -showversion	This option specifies the JDK version. This option exists for JDK 1.3 and 1.4. «
» -verbose[:class gcl jni] «	A message is displayed for each garbage collection sweep.
» -Xrun[:]	A message is displayed, indicating that a service program and an optional parameter string for the JVM_OnLoad function during JVM startup. «



The Run Java (RUNJVA) command in the CL command reference information describes these new options in detail. The CL command reference information for the Create Java Program (CRTJVAPGM) command, Delete Java Program (DLTJVAPGM) command, and Display Java Program (DSPJVAPGM) command contains information about managing Java programs.

The java command in Qshell is available using the Qshell Interpreter.

For more information about the java command in Qshell, see the java tool by Sun Microsystems, Inc.


CL commands that are supported by Java

The IBM Developer Kit for Java^(TM) supports the these CL commands.

-  Analyze Java Virtual Machine (ANZJVM) command retrieves and sets information into a Java virtual machine (JVM). This command helps you debug Java programs by returning information about active classes. 
- Change Java Program (CHGJVAPGM) command changes the attributes of a Java program.
- Create Java Program (CRTJVAPGM) command creates a Java program on an iSeries server from a Java class file, ZIP file, or JAR file.
- Delete Java Program (DLTJVAPGM) command deletes an iSeries Java program that is associated with a Java class file, ZIP file, or JAR file.
- Display Java Program (DSPJVAPGM) command displays information about a Java program on iSeries.
- Dump Java Virtual Machine (DMPJVM) command dumps information about the Java virtual machine for a specified job to a spooled printer file.
- JAVA command and Run Java (RUNJVA) command run iSeries Java programs.

For more information, see the Program and CL Command APIs.

Analyze Java Virtual Machine (ANZJVM) command

 The Analyze Java^(TM) Virtual Machine (ANZJVM) command retrieves and sets information into a Java virtual machine (JVM). The intention is to help you debug Java programs by returning information about active classes.


When you run the ANZJVM command, there is a parameter to specify if a garbage collection cycle should be forced, and if so, an attempt is made to force a garbage collection cycle before each pass. If no garbage collection cycles for the JVM being analyzed have been run, the cycles cannot be forced. There are also parameters on how the information should be stored and how long the interval between passes should be.

After the ANZJVM command completes, a spooled output file is produced. Refer to the Example: ANZJVM command and the output file for more information.

Refer to the ANZJVM command and ANZJVM syntax diagram for more information.

Refer to Considerations for ANZJVM command for more information. 

Run the ANZJVM command

 You can run the Analyze Java^(TM) Virtual Machine (ANZJVM) command to collect information about the Java virtual machine (JVM) for a specified job. By taking a copy of the JVM and comparing the data with another copy taken at a later time, the data can be analyzed to help find object leaks. An interval parameter is used to specify the time between heap passes. When the interval is set to zero, there are two passes of the heap, with the second pass starting just after the first pass finishes. Information for both passes is then returned.

The following information is returned for each class in the heap:

1. Class name
2. Garbage collection heap information
 - a. Pass one
 - b. Pass two
 - c. Change in the number of objects in the garbage collection heap

3. Object space used
 - a. Pass one
 - b. Pass two
 - c. Change in object size
4. Global registry information, which is the same information for the global registry as what is listed for the object table.
5. Loader name

Force a garbage collection cycle: To obtain a cleaner view of the heap, it is desirable to look at it as soon as possible after a garbage collection cycle. ANZJVM has a FRCGC parameter to specify if garbage collection should be forced. The possible options are the following:

- *YES
Garbage collection is forced before each ANZJVM sweep of the heap.
- *NO
Garbage collection is not forced by ANZJVM.◀◀

Considerations for ANZJVM command

▶ Due to the length of time ANZJVM can run, it is highly possible that a JVM ends before ANZJVM is able to finish. In the event that the JVM ends, ANZJVM returns the JVAB606 message (that is, JVM ended while processing ANZJVM) along with the data that it was able to obtain.

There is also no upper limit on the number of classes a JVM can handle. If there are more classes than can be handled, ANZJVM should return the data that can be handled along with a message letting you know there was additional information not reported. When the data requires truncating, ANZJVM returns as much information as possible.

The internal parameter is restricted to 3600 seconds (one hour) in length. The number of classes that ANZJVM can return information about is limited by the amount of storage on your system. ◀◀

Example: ANZJVM command

▶ The following ANZJVM command example collects two copies of the JVM 60 seconds apart for a job named QJVACMDSRC for user name JOHN and the job number 099112. The data from the copies is placed in a QSYSPRT printer device file.

```
ANZJVM JOB(099112/JOHN/QJVACMDSRV)
```

An example of the spooled output file follows.

Spooled output file for ANZJVM command: The following is an example of the data contained in a spooled output file after the ANZJVM command is run:

```
Mon Feb 26 15:39:12 CST 2002
Job: 099112/JOHN/QJVACMDSRV
Interval: 10 seconds
Total garbage collection cycles prior to running: 29
Total garbage collection cycles after running: 31
GC forced: NO
.....
. Class loader information .
.....
0 Default class loader
.....
. GC heap information .
.....
Loader
|   Number of pass one objects in the GC heap
|   |   Number of pass two objects in the GC heap
|   |   |   Change in the number of objects in the GC heap
|   |   |   |   Pass one object size (K)
```

					Pass two object size (K)		
					Change in object size (K)		In global registry
							Class name
0	431359	491363	60004	18979	21619	2640	NO java/lang/ String
0	8	8	0	0	0	0	NO sun/misc/ URLClass Path\$ JarLoader
0	4	4	0	0	0	0	NO java/lang/ Object
0	7	7	0	0	0	0	NO java/util/ zip/ Inflater
0	2	2	0	0	0	0	NO java/lang/ Thread
0	2	2	0	0	0	0	NO [Ljava/lang/ Class;
0	13	13	0	0	0	0	NO java/io/ File Descriptor
0	2	2	0	0	0	0	NO java/io/ Buffered Writer
0	4	4	0	0	0	0	NO [Ljava/lang/ String;
0	1	1	0	0	0	0	NO [Ljava/io/ ObjectStream Class\$ ObjectStream ClassEntry;
0	404	404	0	24	24	0	NO [Ljava/util/ HashMap\$ Entry;
0	423	423	0	15	15	0	NO java/util/ jar/ Attributes\$ Name
0	1	1	0	0	0	0	NO java/io/ Buffered InputStream
0	2	2	0	0	0	0	NO java/io/ Buffered Output Stream
0	1	1	0	0	0	0	NO java/security/ Protection Domain
0	1	1	0	0	0	0	NO sun/security/ provider/Sun
0	1	1	0	0	0	0	NO java/io/ File Permission
0	128	128	0	6	6	0	NO java/util/ Hashtable\$ Entry
0	2	2	0	0	0	0	NO java/net/ URLClass Loader\$ ClassFinder
0	1	1	0	0	0	0	NO java/lang/ Runtime
0	1	1	0	0	0	0	NO java/util/ BitSet
0	7	7	0	0	0	0	NO java/util/ jar/

0	2	2	0	0	0	0	NO	JarVerifier
0	22	22	0	1	1	0	NO	java/lang/ ThreadGroup
0	8	8	0	0	0	0	YES	java/util/ Locale
0	37	37	0	125	125	0	YES	java/io/ RandomAccess File
0	1	1	0	0	0	0	NO	[B sun/misc/ Launcher
0	871	871	0	117	117	0	YES	[C sun/misc/ Launcher\$ Factory
0	1	1	0	0	0	0	NO	java/lang/ Class
0	1	1	0	0	0	0	NO	java/util/ Collections\$ EmptyList
0	1	1	0	0	0	0	NO	java/util/ Collections\$ EmptyMap
0	1	1	0	0	0	0	NO	java/lang/ String\$ Case Insensitive Comparator
0	1	1	0	2	2	0	NO	[I java/lang/ OutOf Memory Error
0	3	3	0	0	0	0	NO	[J java/util/ HashMap\$ Entry
0	1	1	0	0	0	0	NO	java/util/ Random
0	800	800	0	41	41	0	NO	java/security/ Access Control Context
0	5	5	0	0	0	0	YES	java/lang/ ref/ Reference\$ Reference Handler
0	1	1	0	8	8	0	NO	[S [Ljava/lang/ ref/Soft Reference;
0	1	1	0	0	0	0	NO	java/io/ Object Stream Class\$ Compare Member ByName
0	7	7	0	0	0	0	NO	java/util/ jar/ JarFile\$ JarFileEntry
0	1	1	0	0	0	0	NO	java/util/ Collections\$ EmptySet
0	1	1	0	0	0	0	NO	[Ljava/

0	1	1	0	0	0	0	NO	security/ cert/ Certificate;
0	1	1	0	0	0	0	NO	java/lang/ ref/ Reference Queue
0	1	1	0	0	0	0	NO	java/util/ Hashtable\$ Empty Enumerator
0	1	1	0	0	0	0	YES	sun/misc/ Launcher\$ AppClass Loader
0	1	1	0	0	0	0	NO	java/lang/ Shutdown\$ Lock
0	17	17	0	0	0	0	NO	java/util/ Vector
0	3	3	0	0	0	0	NO	java/util/ Stack
0	17	17	0	1	1	0	NO	java/net/URL
0	21	21	0	1	1	0	NO	java/lang/ ref/ Finalizer
0	1	1	0	0	0	0	NO	java/io/ Os400 FileSystem
0	1	1	0	0	0	0	NO	java/lang/ Runtime Permission
0	1	1	0	0	0	0	NO	[Ljava/io/ File;
0	2	2	0	0	0	0	NO	sun/io/ CharToByte ISO8859_1
0	1	1	0	0	0	0	NO	sun/misc/ Launcher\$ ExtClass Loader
0	1	1	0	0	0	0	NO	java/lang/ ref/ Soft Reference
0	1	1	0	0	0	0	NO	sun/security/ provider/ PolicyFile
0	1	1	0	0	0	0	NO	java/io/ Object Stream Class\$ Compare Class ByName
0	1	1	0	0	0	0	NO	sun/net/www/ protocol/ file/ Handler
0	8	8	0	0	0	0	NO	java/util/ jar/ JarFile
0	30	30	0	2	2	0	NO	java/util/ Hashtable
0	1	1	0	0	0	0	NO	java/lang/ ref/ Reference\$

0	2	2	0	0	0	0	NO	Lock
0	1	1	0	0	0	0	NO	java/io/ PrintStream
0	4	4	0	0	0	0	NO	java/util/ Hashtable\$ Empty Iterator
0	391	391	0	12	12	0	NO	java/io/File java/util/ jar/ Attributes
0	2	2	0	0	0	0	NO	sun/misc/ URLClassPath
0	2	2	0	0	0	0	NO	java/io/ FileInput Stream
0	2	2	0	0	0	0	NO	java/io/ Output Stream Writer
0	11	11	0	0	0	0	NO	java/util/ ArrayList
0	1	1	0	0	0	0	NO	java/net/ Unknown Content Handler
0	3	3	0	0	0	0	NO	java/lang/ ref/ Reference Queue\$Lock
0	2	2	0	0	0	0	NO	java/io/ FileOutput Stream
0	1	1	0	0	0	0	NO	sun/misc/ URLClass Path\$ FileLoader
0	31	31	0	5246	5246	0	NO	[Ljava/lang/ Object;
0	1	1	0	0	0	0	NO	java/lang/ Class Loader\$ Native Library
0	2	2	0	0	0	0	NO	[Ljava/lang/ Thread;
0	404	404	0	35	35	0	NO	java/util/ HashMap
0	2	2	0	0	0	0	NO	java/lang/ Boolean
0	1	1	0	0	0	0	YES	java/lang/ref/ Finalizer\$ Finalizer Thread
0	1	1	0	0	0	0	NO	sun/security/ provider/ Policy Permissions
0	7	7	0	0	0	0	NO	java/util/jar/ Manifest
0	2	2	0	0	0	0	NO	sun/net/www/ protocol/ jar/ Handler
0	1	1	0	0	0	0	NO	com/sun/ rsajca/ Provider

```

0 1      1      0      0      0      0      NO  java/util/
Collections
$Reverse
Comparator
0 2      2      0      0      0      0      NO  [Ljava/io/
ObjectStream
Field;
0 1      1      0      0      0      0      NO  java/security/
CodeSource
0 34     34     0      5      5      0      NO  [Ljava/util/
Hashtable$
Entry;
0 2      2      0      0      0      0      NO  java/lang/ref/
Reference
Queue$Null
0 2      2      0      0      0      0      NO  java/util/
Properties
0 2      2      0      0      0      0      NO  java/util/
HashSet
0 1      1      0      0      0      0      NO  [Ljava/lang/
ThreadGroup;
0 1      1      0      0      0      0      NO  java/util/
HashMap$
EmptyHash
Iterator
0 7      7      0      0      0      0      NO  java/io/
ByteArray
OutputStream

```

```

.....
. Global registry information
.....

```

```

Loader
Number of pass one objects in the GC heap
Number of pass two objects in the GC heap
Change in the number of objects in the GC heap
Pass one object size (K)
Pass two object size (K)
Change in object size (K)
Class name
0 8      8      0      0      0      0      java/io/
RandomAccessFile
0 8      8      0      64     64     0      [B
0 10     10     0      1      1      0      [C
0 439   439   0      10     10     0      java/lang/Class
0 1      1      0      0      0      0      java/lang/ref/
Reference$
ReferenceHandler
0 1      1      0      0      0      0      sun/misc/
Launcher$
AppClassLoader
0 1      1      0      0      0      0      java/lang/ref/
Finalizer$
FinalizerThread

```



Example: Change Java Program (CHGJVAPGM) command

To change a Java^(TM) program that is associated with the class file, myJavaClassName, use the Change Java Program (CHGJVAPGM) command. The Java program class file bytecodes are interpreted. To start the program, use the Run Java (RUNJVA) command. The Java program is recreated only if the attributes that are specified are different than those of the current program.

Example 1: Change an interpreted Java program

Note: Read the Code example disclaimer for important legal information.

```
CHGJVAPGM CLSF('/projectA/team2/myJavaClassName.class')
OPTIMIZE(*INTERPRET)
```

The following example is the same as Example 1, except that the program is optimized. It is changed with OPTIMIZE(10), so the program contains compiled machine instructions that run when the Java program starts.

Example 2: Change an optimized Java program in a JAR file

Note: Read the Code example disclaimer for important legal information.

```
CHGJVAPGM CLSF('/projectB/myJavaappfile.jar')
OPTIMIZE(10)
```

See the Change Java Program (CHGJVAPGM) command for syntax diagram and parameter details.

Licensed Internal Code option parameter strings

This table shows the strings that are recognized by the Licensed Internal Code option (LICOPT) parameter. These strings are not case sensitive, but they are shown as mixed case for readability.

LICOPT parameter strings

String	Description
AllFieldsVolatile	If set, treats all fields as volatile.
NoAllFieldsVolatile	If set, no fields are treated as volatile.
AllowBindingToLoadedClasses	Indicates that temporary class representations that were created as a result of defineClass calls within a running Java ^(TM) virtual machine may be tightly bound to other class representations within the same Java virtual machine.
NoAllowBindingToLoadedClasses	Indicates that temporary class representations that were created as a result of defineClass calls within a running Java virtual machine may not be tightly bound to other class representations within the same Java virtual machine.
AllowClassCloning	When multiple Java programs are generated for a JAR file, allows copies of classes from one program to be included in the generated code for another program. Facilitates aggressive inlining.
NoAllowClassCloning	Does not allow copies of classes from one program to be included in the generated code for another program.
AllowInterJarBinding	Allows tight binding to classes outside the class or JAR file being compiled. Facilitates aggressive optimizations.
NoAllowInterJarBinding	Does not allow tight binding to classes outside the class or JAR file being compiled. This overrides the presence of the CLASSPATH and JDKVER parameters on CRTJVAPGM.
➤ AllowMultiThreadedCreate	Indicates that CRTJVAPGM performs as usual, using only one thread. ⏪
➤ NoAllowMultiThreadedCreate	CRTJVAPGM uses multiple threads, if they are available, during creation. ⏪

String	Description
AnalyzeObjectLifetimes	Performs analysis using visible classes to determine which objects are short-lived. A short-lived object does not outlive the method in which it is allocated, and may be subject to more aggressive optimizations.
NoAnalyzeObjectLifetimes	Does not perform analysis of short-lived objects.
AllowBindingWithinJar	Indicates that class representations within a ZIP file or JAR file may be tightly bound to other class representations within the same ZIP file or JAR file.
NoAllowBindingWithinJar	Indicates that class representations within a ZIP file or JAR file may not be tightly bound to other class representations within the same ZIP file or JAR file.
AllowInlining	Tells the translator that it is permitted to inline local methods. This is the default for optimization levels 30 and 40.
NoAllowInlining	Does not tell the translator that it is permitted to inline local methods.
AssumeUnknownFieldsNonvolatile	When the attributes of a field in an external class cannot be determined, this parameter generates code by assuming that the field is non-volatile.
NoAssumeUnknownFieldsNonvolatile	When the attributes of a field in an external class cannot be determined, this parameter generates code by assuming that the field is volatile.
BindErrorHandling	Specifies what action should be taken if, as a result of honoring the AssumeUnknownFieldsNonvolatile, PreresolveExtRef, or PreLoadExtRef Licensed Internal Code option, the Java virtual machine class loader detects that a class representation contains method representations, which cannot be used in the current context.
BindInit	Uses bound call to local init methods.
NoBindInit	Does not use bound call to local init methods.
BindSpecial	Uses bound call to local special methods.
NoBindSpecial	Does not use bound call to local special methods.
BindStatic	Uses bound call to local static methods.
NoBindStatic	Does not use bound call to local static methods.
BindTrivialFields	Binds trivial field references during program creation.
NoBindTrivialFields	Resolves field references at first touch.
BindVirtual	Uses bound call to local final virtual methods.
NoBindVirtual	Does not use bound call to local final virtual methods.
DeferResolveOnClass	Takes a string parameter that is presumed to be the name of a class (for example, java.lang.Integer). When you set PreresolveExtRef to optimization level 40, classes that are specified with DeferResolveOnClass are not in the preresolve operation. This is useful if some classes in unused paths in the code are not in the CLASSPATH. It allows you to use optimization level 40 regardless of this by specifying a "DeferResolveOnClass='somepath.someclass'" for each missing class. Multiple DeferResolveOnClass entries are allowed.

String	Description
DevirtualizeFinalJDK	Allows CRTJVAPGM to use knowledge of the standard JDK to devirtualize calls to those JDK methods that are known to be final methods or members of final classes. It is the default at optimization levels 30 and 40.
NoDevirtualizeFinalJDK	Does not allow CRTJVAPGM to use knowledge of the standard JDK to devirtualize calls to those JDK methods that are known to be final methods or members of final classes.
DevirtualizeRecursive	Causes special code to be generated in the case of some recursive methods and eliminates much of the overhead of the recursive method calls. However, additional checking logic is generated on initial entry to the recursive method, so performance may not improve in cases of shallow recursion.
NoDevirtualizeRecursive	Does not cause special code to be generated in the case of some recursive methods.
DisableIntCse	Causes certain common subexpression optimizations to be disabled when generating code for certain types of integer expressions. This may improve overall optimization by exposing other optimization opportunities to the Optimizing Translator.
NoDisableIntCse	Causes certain common subexpression optimizations to not be disabled when generating code for certain types of integer expressions. This generally results in better performing code at lower optimization levels.
DoExtBlockCSE	Performs extended basic block common subexpression elimination.
NoDoExtBlockCSE	Does not perform extended basic block common subexpression elimination.
DoLocalCSE	Performs local common subexpression elimination.
NoDoLocalCSE	Does not perform local common subexpression elimination.
EnableCseForCastCheck	If set, generates code for castcheck that can be DAGed to an earlier instance.
NoEnableCseForCastCheck	Is not set; does not generate code for castcheck that can be DAGed to an earlier instance.
ErrorReporting	Runtime error reporting field**: Provides the option to fail the compile when encountering verification or class format errors. 0=Report all errors immediately; 0=Report all errors immediately; 1=Defer reporting of bytecode verification errors; 2=Defer reporting of bytecode verification errors and class format errors to runtime.
HideInternalMethods	Causes methods in cloned classes to be made internal, allowing the methods to be omitted if there are no references to them or if all references are inlined. The default is HideInternalMethods for optimization 40 and NoHideInternalMethods for optimization between 0 and 30. <<
InlineArrayCopy	Causes the inlining of the System.arraycopy method in some cases of scalar arrays.
NoInlineArrayCopy	Prevents the inlining of the System.arraycopy method.

String	Description
InlineInit	Inlines init methods for java.lang classes.
NoInlineInit	Does not inline init methods.
InlineMiscFloat	Inlines miscellaneous float/double methods from java.lang.Math.
NoInlineMiscFloat	Does not inline miscellaneous float/double methods.
InlineMiscInt	Inlines miscellaneous int/long methods from java.lang.Math.
NoInlineMiscInt	Does not inline miscellaneous int/long methods.
InlineStringMethods	Permits inlining of certain methods from java/lang/String.
NoInlineStringMethods	Inhibits inlining of certain methods from java/lang/String.
InlineTransFloat	Inlines transcendental float/double methods from java.lang.Math.
NoInlineTransFloat	Does not inline transcendental float/double methods.
OptimizeJsr	Generates better code for "jsr" bytecodes that have a single target.
NoOptimizeJsr	Suppresses generation of better code for "jsr" bytecodes that have a single target.
PreloadExtRef	Indicates that referenced classes may be preloaded (without class initialization) upon method entry.
NoPreloadExtRef	Indicates that referenced classes may not be preloaded upon method entry. However, the PreresolveExtRef parameter overrides this setting and causes referenced classes to be preloaded and initialized.
PreresolveExtRef	Preresolves referenced methods at method entry.
NoPreresolveExtRef	Resolves method references at first touch. Use to resolve "class not found" exceptions on programs that run on other machines.
ProgramSizeFactor	When a JAR file may be large enough to require multiple Java programs, this numeric value (default 100) is used to determine how large each program can grow.
ShortCktAthrow	If set, attempt to short-circuit athrows.
NoShortCktAthrow	Is not set, does not attempt to short-circuit athrows.
ShortCktExSubclasses	If set, recognizes some subclasses of Exception and short-circuit them directly.
NoShortCktExSubclasses	If not set, does not recognize some subclasses of Exception and short-circuit them directly.
StrictFloat	Inhibits floating-point optimizations that are not strictly compliant with the Java specification.
NoStrictFloat	Permits floating-point optimizations that are not strictly compliant with the Java specification.

The double asterisk (**) signifies that these strings require a numerical value for input in the syntax of stringname=number (with no spaces in between).

Example: Create Java Program (CRTJVAPGM) command

To create a Java^(TM) program and associate it with the class file, myJavaClassName, use the Create Java Program (CRTJVAPGM) command. When created with OPTIMIZE(*INTERPRET), the Java program class file bytecodes are interpreted. To start the program, use the Run Java (RUNJVA) command.

Example 1: Create an interpreted Java program

Note: Read the Code example disclaimer for important legal information.

```
CRTJVAPGM CLSF('/projectA/team2/myJavaClassName.class')
OPTIMIZE(*INTERPRET)
```

This example is the same as Example 1, except that the program is optimized. It is created with OPTIMIZE(40), so the program contains compiled machine instructions that run when the Java program starts.

Example 2: Create an optimized Java program

Note: Read the Code example disclaimer for important legal information.

```
CRTJVAPGM CLSF('/projectB/team2/myJavaClassName.class')
OPTIMIZE(40)
```

See the Create Java Program (CRTJVAPGM) command for syntax diagram and parameter details.

Example: Delete Java Program (DLTJVAPGM) command

The Delete Java^(TM) Program (DLTJVAPGM) command deletes the Java program that is associated with the specified class file that is named myJavaClassName.

Note: The DLTJVAPGM command does not delete the class file or ZIP file.

Example 1: Delete a Java program

Note: Read the Code example disclaimer for important legal information.

```
DLTJVAPGM CLSF('/projectA/team2/myJavaClassName.class')
```

See the Delete Java Program (DLTJVAPGM) command for syntax diagram and parameter details.

Example: Dump Java Virtual Machine (DMPJVM) command

The Dump Java^(TM) Virtual Machine (DMPJVM) command dumps information about the Java virtual machine for a specified job.

Example 1: Dump a Java Virtual Machine

Note: Read the Code example disclaimer for important legal information.

```
DMPJVM JOB(099246/FRED/QJVACMSRV)
```

The DMPJVM command dumps the information for the Java virtual machine that is running in the job that is named 099246/FRED/QJVACMSRV.

Example output:

```
JAVA VIRTUAL MACHINE INFORMATION: 099246/FRED/QJVACMSRV
.....
. Classpath
.....
/QIBM/ProdData/Java400/jdk117/lib/jdkptf117.zip:/QIBM/ProdData/Java400/jdk1
```

```
17/lib/classes.zip:/QIBM/ProdData/Java400/ext/IBMmisc.jar:/QIBM/ProdData/Java400/ext/db2_classes.jar:/QIBM/ProdData/Java400/ext/jssl.jar:/QIBM/ProdData/Java400/ext/ibmjssl.jar:/QIBM/ProdData/Java400/./home/fred
```

```
.....
. Garbage collection
.....
Garbage collector parameters
  Initial size: 2048 K
  Max size: *NOMAX
Current values
  Heap size: 9476 K
  Garbage collections: 0
.....
. Thread information
.....
Information for 3 thread(s) of 3 thread(s) processed
Thread: 00000001 Thread-0
  TDE: B000200002941000
  Thread priority: 5
  Thread status: Destroy wait
  Thread group: main
  Runnable: java/lang/Thread
Stack:
  None
Locks:
  None
.....
Thread: 00000003 t2
  TDE: B000100005B37000
  Thread priority: 5
  Thread status: Timed wait
  Thread group: main
  Runnable: dbgtest2
Stack:
  java/io/BufferedInputStream.read()I+11 (BufferedInputStream.java:154)
  pressEnter.theFirstMethod(Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;)V+1
  0 (dbgtest2.java:15)
  dbgtest2.run()V+69 (dbgtest2.java:44)
  java/lang/Thread.run()V+11 (Thread.java:466)
Locks:
  None
.....
Thread: 00000002 t1
  TDE: B000100005B33000
  Thread priority: 5
  Thread status: Java wait
  Thread group: main
  Runnable: dbgtest2
Stack:
  pressEnter.theFirstMethod(Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;Ljava/lang/String;
  Ljava/lang/String;)V+0 (dbgtest2.java:14)
  dbgtest2.run()V+69 (dbgtest2.java:44)
  java/lang/Thread.run()V+11 (Thread.java:466)
Locks:
  None
.....
```

Example: Display Java Program (DSPJVAPGM) command

The DSPJVAPGM (Display Java Program) command displays the Java program that is associated with the specified class file named myJavaClassName.

Example 1: Display a Java Program

Note: Read the Code example disclaimer for important legal information.

```
DSPJVAPGM CLSF('/projectA/team2/myJavaClassName.class') OUTPUT(*)
```

See the Display Java Program (DSPJVAPGM) command for syntax diagram and parameter details.

JAVA command

The JAVA command functions exactly the same as the Run Java^(TM) (RUNJVA) command. You can use them interchangeably. See Run Java (RUNJVA) command for information and parameters that you can use with the JAVA command.

Example: Use the Run Java (RUNJVA) command

The Run Java^(TM) (RUNJVA) command runs the iSeries Java program that is associated with the class.

Example 1: Run a Java program

Note: Read the Code example disclaimer for important legal information.

```
RUNJVA CLASS ('/projectA/myJavaClassName')
```

See the Run Java (RUNJVA) command for syntax diagram and parameter details.

iSeries Navigator commands that are supported by Java

The iSeries Navigator is a graphical interface for your Windows^(R) desktop. It is part of iSeries Access for Windows and covers many iSeries functions that administrators or users need to accomplish their daily work.

iSeries Navigator supports Java^(TM) as a plugin contained in the File Systems option of iSeries Access for Windows. To use the iSeries Navigator Java plugin, you need to install the IBM Developer Kit for Java on your iSeries server. Then, to install the Java plugin on your personal computer, select File Systems through Selective Setup in the Client Access folder.

Class, JAR, ZIP, and Java files reside in the integrated file system. iSeries Navigator allows you to see these files in the right pane. Right-click the class, JAR, ZIP, or java file that you want to use. This brings up a context menu.

➤ Selecting **Associated Java Program** → **New...** from the context menu starts the Java transformer, which creates iSeries Java programs that are associated with your class, JAR, or ZIP file. ⏪ A dialog box allows you to specify details on how to create the program. You can create the programs for either Java transformation or Java interpretation.

Note: If you select transformation, the bytecodes in your class file transform into RISC instructions that result in better performance than if you used interpretation.

➤ Selecting **Associated Java Program** → **Edit...** from the context menu changes attributes of Java programs that are attached to Java class files, ZIP files, or JAR files. ⏪

➤ Selecting **Associated Java Program** → **Run...** from the context menu runs your class file on your iSeries server. ⏪ You may also select a JAR or ZIP file and run a class file located within that JAR or ZIP file. A dialog appears to allow you to specify details on how to run the program. ➤ If you have already selected **Associated Java Program** → **New...**, the iSeries Java program that is associated with your

class file is used when running the program. ⏪ If an iSeries Java program is not already associated with your class file, then the iSeries Java program is created before the program runs.

➤ Selecting **Associated Java Program** → **Delete...** from the context menu deletes the iSeries Java programs that are associated with your class, JAR, or ZIP file. ⏪

Selecting **Properties** from the context menu displays a properties dialog box which contains the **Java Programs** and **Java Options** tabs. These tabs allow you to see the details on how the associated iSeries Java programs were created for your class, JAR, or ZIP file.

Note: These panels are the Display Java Program information.

Selecting **Compile Java file** from the context menu converts any java files that you have selected into their class file bytecodes.

➤ See the help information, included with iSeries Navigator, for the parameters and options of the **New Java Program**, **Edit Java Program**, **Run Java Program**, **Java Programs**, **Java Options**, **Compile Java file**, and **Delete Java Program** iSeries Navigator dialogs. ⏪

Chapter 7. Optional packages

» Optional packages define Application Programming Interfaces (APIs) that extend the core Java platform API. The following are optional packages that you can use with the IBM Developer Kit for Java^(TM):

Java Authentication and Authorization Service

The Java Authentication and Authorization Service (JAAS) allows the association of a specific user or identity to the current Java thread.

Java Cryptography Extension

The Java Cryptography Extension (JCE) provides a framework and implementations for encryption, key generation and key agreement, and Message Authentication Code (MAC) algorithms. JCE also supports secure streams and sealed objects.

Java Naming and Directory Interface

The Java Naming and Directory Interface (JNDI) is part of JavaSoft's platform application program interfaces (API). With JNDI, you can connect to multiple naming and directory services. You can build powerful and portable directory-enabled Java applications by using this interface.

Java Secure Socket Layer

The Java Secure Socket Layer (JSSL) is a set of Java packages that enable secure Internet communications. It implements a Java version of SSL and Transport Layer Security (TLS) protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication.

JavaMail

The JavaMail API provides a set of abstract classes that models an electronic (e-mail) system. The API provides a platform-independent and protocol-independent framework to build Java-based e-mail and messaging applications.

JavaPrintService


The Java Print Service API allows printing on all Java platforms. Java 1.4 provides a framework in which Java runtime environments and third parties can provide steam generator plugins for producing various formats for printing, such as PDF, Postscript, and Advanced Function Presentation^(TM) (AFP^(TM)).




Java Naming and Directory Interface


The Java^(TM) Naming and Directory Interface (JNDI) is part of JavaSoft's platform application program interfaces (API). With JNDI, you can connect seamlessly to multiple naming and directory services. You can build powerful and portable directory-enabled Java applications by using this interface.

JavaSoft developed the JNDI specification with leading industry partners, including IBM, SunSoft, Novell, Netscape, and Hewlett-Packard Co.

For more information about JNDI, see [Java Naming and Directory interface by Sun Microsystems, Inc.](#)  For IBM-specific information, see the [IBM JNDI LDAP provider programming guide](#).

IBM JNDI LDAP provider programming guide

➤ This programming guide assumes that you are familiar with the Java^(TM) Naming and Directory Interface (JNDI) and how Lightweight Directory Access Protocol (LDAP) operates. For more information, see the JNDI documentation  from Sun Microsystems.

IBM provides a LDAP service provider for JNDI that can be used with SDK or Java Runtime Environment (JRE) 1.2.2. Using the IBM JNDI LDAP provider with SDK or JRE 1.3 or later is not supported; instead, JNDI and the Sun Microsystems, Inc. JNDI LDAP provider should be used and are part of the SDK and JRE 1.3. The Sun Microsystems, Inc. JNDI LDAP provider can also be used with SDK and JRE 1.2.2, but you must download these components from the Sun Microsystems, Inc. JNDI  Web site, and support is provided by Sun. This programming guide describes using the IBM JNDI LDAP provider with SDK or JRE 1.2.2.

To compile or run code using the IBM JNDI LDAP provider, add the following to your classpath:

```
/QIBM/ProdData/Java400/ext/ibmjndi.jar:/QIBM/ProdData/Java400/ext/jndi.jar
```

This programming guide discusses the following topics:

Create an initial context

This topic describes how to create an initial context to connect with an LDAP server. JNDI supports two ways for a client to work with a Lightweight Directory Access Protocol (LDAP) server:

- The client identifies the server at context creation.
- A URL string is directly passed to the context's methods.

LDAP V3 URL

This topic defines the LDAP URL syntax.

Server binding and SASL support

A server must authenticate a client before certain operations are allowed. LDAP refers to this as binding to the server. The LDAP protocol extended its authentication to also support Simple Authentication and Security Layer (SASL) mechanisms. These mechanisms allow more sophisticated ways of identifying a client to the server without unnecessarily compromising the security of the user by sending his ID and password in clear text.

Search and obtain attributes

JNDI provides flexibility in searching Lightweight Directory Access Protocol (LDAP) directories.

Add and delete entries in a directory

JNDI allows you to add and delete entries in a directory. This topic includes examples on how to perform these tasks.

Change attributes

JNDI makes it possible to change, create, or remove attributes from a directory entry.

Rename a directory entry

JNDI allows you to rename a directory entry anywhere relative to the base context. This topic identifies a property that affects the rename method used to rename a directory entry.

Referrals and search references

LDAP servers may return either referrals or search references. A referral can be returned on any operation and indicates that the server does not hold the target entry of the request. A search reference is only returned on a search operation.

LDAP controls

The LDAP v3 specification added controls to send and receive extension information. Controls sent to a server are referred to as request controls. Controls received from a server are referred to as response controls.

Binary attributes

The LDAP protocol does not provide any distinction between retrieved binary and textual attributes. Instead, it expects the client application to know how to handle the data. This topic describes three different methods for handling what happens to retrieved attributes and whether or not they are converted to strings.

Schema

You can retrieve, view, and update an LDAP server's schema structure. Only servers that provide schema information defined by the Lightweight Directory Access Protocol specification are supported.

SASL plug-in

You can write your own Simple Authentication and Security Layer (SASL) plug-in. This topic provides code examples to help you get started with creating the plug-in.

Client-side caching

Caching provides a means to store recently requested information locally. This improves performance by retrieving repeated queries locally instead of going back to a remote server for information already obtained.

Retrieve the IBMJNDI class version

This topic indicates what static method you can use to retrieve the IBMJNDI class version of LDAP.

Conformance considerations and additional properties

See this topic to know what considerations must be taken when using the IBM LDAP provider and Sun's JNDI Implementor Guidelines for LDAP Service Providers. Properties that are supported and have been deprecated are also indicated.



Create an initial context

» The Java^(TM) Naming and Directory Interface (JNDI) supports two different ways for a client to work with a Lightweight Directory Access Protocol (LDAP) server. The first and most common way is for the client to identify the server at context creation. Operations are then performed on this open connection by passing DN-based names to the context's methods. The following two properties support this type of operation:

java.naming.factory.initial (Context.INITIAL_CONTEXT_FACTORY)

This property must be set to `com.ibm.jndi.LDAPCtxFactory`.

java.naming.provider.url (Context.PROVIDER_URL)

This property identifies the name and port of the LDAP server in the form of a URL string. If the LDAP server's name and port cannot be identified, the IBM JNDI LDAP provider defaults to `ldap://localhost:389`.

The following code creates a connection to host `ldapservers` and retrieves an entry:

Example 1: Create a connection to host `ldapservers`

Note: Read the Code example disclaimer for important legal information.

```

Properties env = new Properties();
env.put("java.naming.factory.initial", "com.ibm.jndi.LDAPCtxFactory");
env.put("java.naming.provider.url", "ldap://ldapsrvr");
DirContext ctx = new InitialDirContext(env);
Attributes entry = ctx.getAttributes("cn=example,o=IBM,c=US");

```

The second way to work with an LDAP server using JNDI is to pass a URL string directly to the context's methods. This process, however, has the overhead of creating a new connection for each operation and should be avoided if all operations are bound for a single server. The following property supports this type of operation:

java.naming.factory.url.pkgs (Context.URL_PKG_PREFIXES)

This property must be set to `com.ibm.jndi` if passing URL strings as name input to a context's method and connecting to an LDAP server is not required at context creation time.

The following code duplicates the previous example, but delays connecting to the server until the `getAttributes` method is called:

Example 2: Create a connection to host `ldapsrvr` and delay connecting to the server

Note: Read the Code example disclaimer for important legal information.

```


Properties env = new Properties();
env.put("java.naming.factory.url.pkgs", "com.ibm.jndi");
DirContext ctx = new InitialDirContext(env);
Attributes entry = ctx.getAttributes("ldap://ldapsrvr/cn=example,o=IBM,c=US");

```

The provider also supports mixing the two previous methods. That is, it is possible to establish a connection to an LDAP server using `java.naming.factory.initial` and then pass a URL string as a name input to a method. This works whether `java.naming.factory.url.pkgs` is defined or not.

Note: Any connection to the LDAP server opened by **InitialDirContext** must be closed by calling the `close` method. <<

LDAP V3 URL

>> The IBM Java[™] Naming and Directory Interface (JNDI) Lightweight Directory Access Protocol (LDAP) provider fully supports the LDAP Uniform Resource Locator (URL) format defined in RFC 2255 . An LDAP URL is defined by the following syntax:

```

scheme "://" [host [ ":" port ]] [ "/"
  [dn ["?" [attributes] ["?" [scope]
    ["?" [filter] ["?" extensions]]]]]]

```

Where:

- *scheme* indicates the URL scheme. This class library supports either the traditional `ldap` for normal LDAP connections or `ldaps` for Secure Socket Layer (SSL) connections.
- *host* is the name of the LDAP server. If the LDAP server name is not specified, the default is `localhost`.
- *port* represents the LDAP server's port number. If the port number is not specified, the default is 389 for non-SSL and 636 for SSL.
- *dn* identifies the base object for the operation.
- *attributes* represents a comma-separated list of attributes to be returned. If a list of attributes is not specified, the default is to return all attributes.
- *scope* represents the scope of a search. Valid values for these files are the following:
 - **base**
This value represents the base object.
 - **sub**
This value represents the suborder of a file hierarchy.

- **one**
This value represents one level of the file hierarchy.

If the scope is not specified, the default is base.

- *filter* represents the search filter. If the search filter not specified, the default is (objectclass=*).
- *extension* provides the LDAP URL with an extensibility mechanism, allowing the capabilities of the URL to be extended. The only extension supported by the IBM JNDI LDAP provider is bindname.

Server binding and SASL support

➤ In many cases, a server must authenticate a client before certain operations are allowed. The Lightweight Directory Access Protocol (LDAP) refers to this as binding to the server.

When binding to the server, the client specifies which LDAP protocol it wants to use. There are two versions of the LDAP protocol defined, V2 and V3. If the server only supports V2, a protocol error is returned when a client attempts to bind as a V3 client. The IBM Java^(TM) Naming and Directory Interface (JNDI) LDAP provider supports binding as either a V2 or V3 client.

You can use the following properties when binding to a server:

java.naming.ldap.version

This property specifies the LDAP protocol version. Valid values are 2 or 3. If this property is not set, the provider attempts to bind as a V3 client and then automatically steps down to V2 if a protocol error is returned. The provider does not attempt to step down if this property is set.

Besides setting a protocol version, a bind also identifies a user to the server for authentication.

java.naming.security.principal (Context.SECURITY_PRINCIPAL)

This property specifies the client's ID. In almost all cases, it is in the form of a distinguished name.

java.naming.security.credentials (Context.SECURITY_CREDENTIALS)

This property specifies the client's credentials (that is, the client's password).

LDAP also supports different types of authentication mechanisms. The V2 LDAP protocol only supported one type of bind that is referred to as a simple bind. With this mechanism, the clear text ID and credentials are sent to the server. The V3 protocol extended authentication to also support Simple Authentication and Security Layer (SASL) mechanisms. These mechanisms allow more sophisticated ways of identifying a client to the server without unnecessarily compromising the security of the user by sending his ID and password in clear text.

The provider supports two different ways of specifying the authentication mechanism. One way requires the name of the authentication class. This way allows the provider to be extended by specifying authentication classes outside the provider. Thus, you can write your own SASL plug-ins. The following properties support this way of specifying the authentication mechanism.

java.naming.security.sasl

This property specifies the name of the authentication class to use. The following classes are shipped as part of the provider.

com.ibm.ldap.LDAPSimpleBind

This property specifies the clear text ID and credentials that are sent to the server for authentication. This mechanism is supported by both V2 and V3 servers. Note that the use of clear text passwords is not recommended over open networks when there is no authentication or encryption being performed by a lower layer.

com.ibm.ldap.LDAPSaslExternal

The external SASL method attempts to bind using the underlying security protocol already negotiated such as SSL. In most cases, the security principal and credentials should be left uninitialized.

com.ibm.ldap.LDAPSaslCRAM_MD5

The CRAM-MD5 SASL sends the security principal and credentials to the server for authentication using a challenge-response protocol.

com.ibm.ldap.LDAPSaslGSSAPI

The GSSAPI SASL method attempts to bind using Kerberos authentication after obtaining credentials through a separate means such as kinit or integrated login. In most cases, the security principal and credentials bind arguments should be left uninitialized.

java.naming.sasl.mode

This property specifies a mode setting that is passed to the loaded SASL plug-in. This setting is ignored by all predefined SASL plug-ins in the provider.

The second supported way to specify the authentication mechanism is compatible with Sun's LDAP toolkit. Instead of indicating the authentication class to load, the name of the authentication mechanism is specified instead. The IBM JNDI LDAP provider uses this method for specifying the authentication mechanism if the `java.naming.security.sasl` property is not set.

java.naming.security.authentication (Context.SECURITY_AUTHENTICATION)

This property specifies the name of the authentication mechanism to use. The following values are supported by this property.

- none
Do not perform authentication (anonymous bind).
- simple
Use simple authentication.
- EXTERNAL
Use external SASL mechanism.
- CRAM-MD5
Use CRAM-MD5 SASL mechanism.
- GSSAPI
Use GSS or Kerberos SASL mechanism.

The following property is set by this class library after the client has been successfully authenticated:

java.naming.authorization.identity

This property is set to the authorization identity of the client. Normally, this is the same as the specified client ID. However, the SASL mechanism may map the initial bind DN to another value. An example of this is the external SASL where the authorization identity is stored in the client's certificate.

The following example demonstrates setting properties to indicate that the version protocol to 3 and authenticating as Larry Meade using the CRAM-MD5 mechanism:

Example: Set properties

Note: Read the Code example disclaimer for important legal information.

```
Properties env = new Properties();
env.put("java.naming.factory.initial", "com.ibm.jndi.LDAPCtxFactory");
env.put("java.naming.ldap.version", "3");
env.put("java.naming.provider.url", "ldap://ldapservers");
env.put(Context.SECURITY_PRINCIPAL, "cn=Larry Meade, o=IBM, c=US");
env.put(Context.SECURITY_CREDENTIALS, "secret");
env.put(Context.SECURITY_AUTHENTICATION, "CRAM-MD5");
DirContext ctx = new InitialDirContext(env);
```

The previous example can be changed to explicitly indicate the SASL class name. The `SECURITY_AUTHENTICATION` line in the previous example would need to be replaced with the following line:

```
env.put("java.naming.security.sasl", "com.ibm.ldap.LDAPSaslCRAM_MD5");
```

Search and obtain attributes

» The Java^(TM) Naming and Directory Interface (JNDI) provides great flexibility in searching Lightweight Directory Access Protocol (LDAP) directories. The two most frequently used methods in the IBM JNDI LDAP provider are `search` and `getAttributes`. However, the following methods also retrieve data from an LDAP server:

- `lookup`
- `lookupLink`
- `list`
- `listBindings`
- `getSchema`
- `getSchemaClassDefinition`

The following properties affect search operations:

java.naming.ldap.derefAliases

This property defines how alias objects (as defined in X.501) are to be handled. This property accepts the following values:

- `always`
This value dereferences aliases when searching and locating the base object of the search. This is the default.
- `never`
This value does not dereference aliases when searching or in locating the base object of the search. To increase performance, this is the recommended setting.
- `finding`
This value dereferences aliases when locating the base object of the search, but not when searching subordinates of the base object.
- `searching`
This value dereferences aliases in subordinates of the base object in searching, but not in locating the base object of the search.

java.naming.batchsize (Context.BATCHSIZE)

This value sets the recommended size limit for the number of search results held by a returned `NamingEnumeration`. If the value is not specified, then the default batchsize is 1. This helps to ensure the smallest memory footprint possible by the class library. A value of 0 disables batchsize and indicates that a search is blocked until all results are collected.

java.naming.ldap.typesOnly

This property is relevant to the `getAttributes` and `search` methods, the latter only when the returning object flag is false. This property accepts the following values:

- `true`
This value returns only attribute identifiers, but not the values.
- `false`
This value returns both the attribute identifiers and values. This is the default.

The result of a search call is a `NamingEnumeration`. To obtain the results, the enumeration must be walked using either the traditional `hasMoreElements` and `nextElement` methods or the `NamingEnumeration`-specific `hasMore` and `next` methods. The latter two methods allow exceptions to be caught if you want to chase or view a `ReferralException`.

Notes:

- To avoid unexpectedly keeping resources allocated and connections open, a NamingEnumeration should either be walked to the end (that is, until the hasMore and hasMoreElements methods return false) or the enumeration's close method should be called.
- The IBM JNDI LDAP provider automatically replaces invalid UTF-8 character encodings in distinguished names and attribute type names with a * character. This is done to prevent a single invalid value from causing an entire and possibly lengthy search to fail.

The following example performs a search for all entries where the surname is smith and indicates return of only the cn attribute:

Example: Search and obtain the cn attribute

Note: Read the Code example disclaimer for important legal information.

```
SearchControls constraints = new SearchControls();
constraints.setSearchScope(SearchControls.SUBTREE_SCOPE);
String attrList[] = {"cn"};
constraints.setReturningAttributes(attrList);
NamingEnumeration results =
    ctx.search("o=IBM,c=US", "(sn=smith)", constraints);
while (results.hasMore()) {
    SearchResult si =(SearchResult)results.next();
    System.out.println(si.getName());
    Attributes attrs = si.getAttributes();
    if (attrs == null) {
        System.out.println("  No attributes");
        continue;
    }
    NamingEnumeration ae = attrs.getAll();
    while (ae.hasMoreElements()) {
        Attribute attr =(Attribute)ae.next();
        String id = attr.getID();
        Enumeration vals = attr.getAll();
        while (vals.hasMoreElements())
            System.out.println("  "+id + ": " + vals.nextElement());
    }
}
```

The following example uses the list method to display the names under a base distinguished name (DN):

Example: Display names under a base DN

Note: Read the Code example disclaimer for important legal information.

```
String url="ldap://ldapsrv:389/o=IBM,c=US";
NamingEnumeration listResults=ctx.list(url);
while (listResults.hasMore()) {
    NameClassPair ncp = (NameClassPair) listResults.next();
    System.out.println(ncp.getName());
}
```



Add and delete entries in a directory

➤ The Java^(TM) Naming and Directory Interface (JNDI) allows you to add and delete entries in a directory. The following example adds a new entry with objectclass, roomnumber, and telephonenumber attributes:

Example: Add an entry in a directory

Note: Read the Code example disclaimer for important legal information.


```

BasicAttribute objClasses = new BasicAttribute("objectclass");
objClasses.add("person");
objClasses.add("organizationalPerson");
objClasses.add("inetOrgPerson");

BasicAttributes attrs = new BasicAttributes();
attrs.put(objClasses);
attrs.put("roomnumber", "2000");
attrs.put("telephonenumber", "1-800-use-LDAP");

ctx.createSubcontext(name, attrs);

```

The following example removes an entry:

Example: Delete an entry in a directory

Note: Read the Code example disclaimer for important legal information.

```
ctx.destroySubcontext(name);
```

Change attributes: JNDI makes it possible to change, create, or remove attributes from a directory entry. The following example replaces the roomnumber attribute of an entry:

Example: Change an attribute

Note: Read the Code example disclaimer for important legal information.

```
ctx.modifyAttributes(name,
    DirContext.REPLACE_ATTRIBUTE,
    new BasicAttributes("roomnumber", "5000"));
```

The following example adds a new telephonenumber attribute value to an entry and removes the roomnumber attribute:

Example: Change an entry from a directory

Note: Read the Code example disclaimer for important legal information.

```
ModificationItem[] mods=new ModificationItem[2];
mods[0] = new ModificationItem(DirContext.ADD_ATTRIBUTE,
    new BasicAttribute("telephonenumber", "456-7777"));
mods[1] = new ModificationItem(DirContext.REMOVE_ATTRIBUTE,
    new BasicAttribute("roomnumber"));
ctx.modifyAttributes(name, mods);
```

Rename a directory entry: You can use the rename method to rename a directory entry anywhere relative to the base context.

The following property affects the rename method:

java.naming ldap.deleteRDN

This property removes the old RDN when the entry is renamed. The default setting is true.

If this property is set to false, the old RDN is retained as an attribute value of the entry.

The following is an example of calling the rename method:

Example: Rename a directory entry

Note: Read the Code example disclaimer for important legal information.

```
String oldname="cn=bill smith";
String newname="cn=bill smith, ou=programmer";
ctx.rename(oldname, newname);
```



Referrals and search references

▶ Lightweight Directory Access Protocol (LDAP) servers may return either referrals or search references. A referral can be returned on any operation and indicates that the server does not hold the target entry of the request. A search reference is only returned on a search operation. The search reference indicates that the server was able to locate the entry referred to by the baseObject, but was unable to search all the entries in the scope at and under the baseObject. A server may return one or more search references.

A context may be configured to handle referrals and search references in one of three ways:

1. It can be set to automatically follow the reference and perform the operation at the referred to server. The IBM Java^(TM) Naming and Directory Interface (JNDI) LDAP provider automatically recognizes and avoids referral loops; that is, situations where a referral points back to one already chased earlier in the chain.
2. It can be set to throw a `ReferralException` when either a referral or search reference is received. This is useful if the automated handling is somehow deficient; for example, each server requires a different binding.
3. It can be set to ignore the reference and continue as if nothing happened. In the case of a search reference, this would mean that only entries found at the originating server are returned.

The following environment properties are defined by the IBM JNDI LDAP provider to handle referrals and search references:

java.naming.referral (Context.REFERRAL)

This property is set to either follow, throw, or ignore. If the property is not set, the default is set to automatically follow referrals.

java.naming.ldap.referral.limit

This property defines the number of referral hops the class library makes when chasing referrals. If the property value is not specified, the default is 10.


java.naming.ldap.referral.bind

If this property is set to true, the classes, when automatically following referrals, bind to any referred to server using the same SASL mechanism and credentials as with the originating context. If set to false, the classes do not bind (that is, anonymous access). The default behavior is to bind.

The following rules apply when chasing referrals or search references:

- If the reference contains a port, then it is used. Otherwise, the port from the primary connection is used.
- The type of security connection from the primary connection is maintained. That is, if the primary connection was over SSL, then all chased referrals are also over SSL.

Note: A context may still throw a `ReferralException` even when set to follow. This can occur if the referral hop limit is exceeded or if the context cannot connect or bind to any of the referred to servers.

The following example catches and displays a referral and search references on a search request. For more information on handling referral exceptions, see `ReferralException`  by Sun Microsystems, Inc.

Example: Catch and display a referral and search references

Note: Read the Code example disclaimer for important legal information.

```
ctx.addToEnvironment(ctx.REFERRAL, "throw");
try {
    NamingEnumeration results = ctx.search(url);
    while (true) {
        try {
```

```

        if (!results.hasMore())
            break;
        SearchResult si =(SearchResult) results.next();
        System.out.println(si.getName());
    } catch (ReferralException re) {
        System.out.println("Reference caught");
        do {
            System.out.println(re.getReferralInfo());
        } while (re.skipReferral());
    }
}
} catch (ReferralException re) {
    System.out.println("Referral caught");
    do {
        System.out.println(re.getReferralInfo());
    } while (re.skipReferral());
}
}

```

LDAP controls

» The Lightweight Directory Access Protocol (LDAP) v3 specification added controls to send and receive extension information. Controls sent to a server are referred to as **request controls**, an example of which would be one which tells the server to sort the results of a search by a specified attribute. What request controls are supported is completely dependent on the server (that is, a control may work on one type of server, but fail on another). Controls received from a server are referred to as **response controls**.

The Java^(TM) Naming and Directory Interface (JNDI) 1.2 specification separates request controls into two distinct categories, those used when connecting to a server and those used on any other operation. For more information about request controls, refer to `LdapContext` in Sun's JNDI documentation.

The IBM JNDI LDAP provider comes with one predefined control called `ManageDsaIT`. This control forces the server to treat search references as normal LDAP entries, allowing them to be viewed and changed instead of the data to which they refer. The following example demonstrates how to enable this control:

Example: Enable the `ManageDsaIT` control

Note: Read the Code example disclaimer for important legal information.

```

import com.ibm.jndi.ldap.control.ManageDsaIT;

Control[] cntl = new Control[1];
cntl[0] = new ManageDsaIT();
ctx.setRequestControls(cntl);

```

Connect controls are only active when the client is binding to a server.

The following is an example of how to enable connect controls:

```

LdapContext ctx = new InitialLdapContext(env, cntl);

```

The following API is used to retrieve the last received response controls:

```

Control[] cntl = ctx.getResponseControls();

```

Refer to Sun's documentation for details on enabling control factories capable of mapping raw control data to specific control classes. <<

Binary attributes

» The Lightweight Directory Access Protocol (LDAP) protocol does not provide any distinction between retrieved binary and textual attributes. Instead, it expects the client application to know how to handle the data. The IBM JNDI LDAP provider is helpful in that it converts textual attributes and returns them as Java strings. However, the provider needs to know which attributes are binary and which attributes represent

character data. The provider supports three different methods for handling what happens to retrieved attributes and whether or not they are converted to strings.

When an attribute is retrieved, the provider checks a list of known binary attribute names. The provider is programmed to recognize the following set of common LDAP binary attributes:

- userPassword
- userCertificate
- cACertificate
- authorityRevocationList
- certificateRevocationList
- deltaRevocationList
- crossCertificatePair
- x500UniqueIdentifier
- photo
- personalSignature
- audio
- jpegPhoto
- javaSerializedObject
- thumbnailPhoto
- thumbnailLogo
- supportedAlgorithms
- protocolInformation

You may specify their own list of binary attribute names using the following property:

java.naming.ldap.attributes.binary (LDAPCtx.ATTRIBUTES_BINARY)

A space-separated list of user-defined binary attribute names in addition to the default set defined by the provider.

The following example identifies two additional user-defined binary attributes:

Example: Specify a list of binary attribute names


Note: Read the Code example disclaimer for important legal information.

```
ctx.addToEnvironment(LDAPCtx.ATTRIBUTES_BINARY,  
    "gifPhoto fingerPrint");
```

The second method for handling binary attributes is recognition of the binary description option supported by some V3 servers. When ";binary" is appended to the attribute name such as "jpegPhoto;binary", it is designated a binary value. The provider recognizes a binary attribute description option when bound as a V3 user.


Finally, the provider attempts to convert any attribute which is not defined as binary and does not contain the binary attribute description option. If the conversion fails, the data is returned as binary instead. However, you should not rely on this working because it is possible to have non UTF-8 data incorrectly converted. <<

Schema

>> You can retrieve, view, and update a Lightweight Directory Access Protocol (LDAP) server's schema structure. Only servers that provide schema information defined by the Lightweight Directory Access Protocol (v3)  document are supported.

To retrieve a server's schema, use the `getSchema` method. The returned schema is represented as a hierarchy, where each subordinate level is a different component of the schema. The IBM Java^(TM) Naming and Directory Interface (JNDI) LDAP provider is capable of parsing the following schema components:

- AttributeTypes, stored under the AttributeDefinition suborder
- Object Classes, stored under the ClassDefinition suborder
- Syntax Descriptions, stored under the SyntaxDefinition suborder
- Matching Rules, stored under the MatchingRule suborder
- IBM Attribute Types, stored under the IBMAttributeDefinition suborder

The content of the entries, excluding IBMAttributeDefinition, has a one-to-one correspondence with the schema defined in Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions  .

IBMAttributeDefinition extends an attribute schema to hold IBM-specific information. It is defined with the following Backus-Naur form (BNF):

```

IBMAttributeTypesDescription = "(" whsp
    numericoid whsp
    [ "DBNAME" qdescrs ] ; at most 2 names (table, column)
    [ "ACCESS-CLASS" whsp IBMAccessClass whsp ]
    [ "LENGTH" wlen whsp ] ; maximum length of attribute
    [ "EQUALITY" [ IBMwlen ] whsp ] ; create index for matching rule
    [ "ORDERING" [ IBMwlen ] whsp ] ; create index for matching rule
    [ "APPROX" [ IBMwlen ] whsp ] ; create index for matching rule
    [ "SUBSTR" [ IBMwlen ] whsp ] ; create index for matching rule
    [ "REVERSE" [ IBMwlen ] whsp ] ; reverse index for substring
    whsp ")"

IBMAccessClass =
    "NORMAL" / ; this is the default
    "SENSITIVE" /
    "CRITICAL" /
    "RESTRICTED" /
    "SYSTEM" /
    "OBJECT"

IBMwlen = whsp len

```

Schema definitions that are returned by the server but are unsupported by this class library are saved and are not parsed. These entries contain exactly two attributes, one with an objectclass equal to the schema type name (for example, objectclass=adddef) and another with the list of values. Unsupported schema definitions may be viewed but not updated.

The following example retrieves the entire schema hierarchy:

Example: Retrieve a schema hierarchy

Note: Read the Code example disclaimer for important legal information.

```

DirContext schemaCtx = ctx.getSchema("");
SearchControls cons = new SearchControls();
cons.setSearchScope(SearchControls.SUBTREE_SCOPE);
NamingEnumeration ne = schemaCtx.search("",
    "(|(NUMERICOID=*)(objectclass=*))", cons);

```

The following example retrieves the schema for the cn attribute:

Example: Retrieve the schema for the cn attribute

Note: Read the Code example disclaimer for important legal information.

```
DirContext schemaCtx = ctx.getSchema("");
Attributes attrs = schemaCtx.getAttributes("AttributeDefinition/cn");
```

The following example attempts to add a schema definition for a new object class:

Example: Add a schema definition

Note: Read the Code example disclaimer for important legal information.

```
DirContext schemaCtx = ctx.getSchema("");
BasicAttributes attrs = new BasicAttributes();
attrs.put("NAME", "javaObject");
attrs.put("NUMERICOID", "1.3.6.1.4.1.42.2.27.4.2.2");
Attribute may = new BasicAttribute("MAY");
may.add("javaClassName");
may.add("javaSerializedObject");
attrs.put(may);
attrs.put("DESC", "Serialized Java object");
attrs.put("AUXILIARY", "true");
attrs.put("SUP", "top");
schemaCtx.createSubcontext("ClassDefinition/javaObject", attrs);
```

The `com.ibm.jndi.LDAPSchemaCtx` class has been specially extended to work with schema definitions from a file. Two public constructors support reading schema information from a disk. One constructor takes as an argument a single filename and the other an array of filenames. The `dumpSchema` method saves the schema definition to a file. The following demonstrates this support:

Example: The `dumpSchema` method

Note: Read the Code example disclaimer for important legal information.

```
LDAPSchemaCtx ctx = new LDAPSchemaCtx("schema.file");
ctx.dumpSchema("schema.sav");
```

SASL plug-in

➤ You can write your own Simple Authentication and Security Layer (SASL) plug-in. The SASL plug-in must be derived from the abstract base class `LDAPSaslBind`. You must implement the `bind` method which then calls the `SendBindRequest` method to communicate with the server. The following is an example of a simple bind plug-in:

Example: Bind plug-in

Note: Read the Code example disclaimer for important legal information.

```
import java.io.IOException;
import com.ibm.asn1.ASN1Exception;
import com.ibm.lldap.*;

public class SimpleBind extends LDAPSaslBind
{
    public boolean bind(String dn, String credentials)
        throws IOException, ASN1Exception, LDAPException
    {
        return SendBindRequest("SIMPLE", dn, credentials);
    }
}
```

If the SASL protocol includes server challenges, the `SendBindRequest` method must be called multiple times, possibly making use of the `getServerCredentials` method to retrieve server challenge information. The following is an example of the CRAM-MD5 SASL plug-in:

Example: CRAM-MD5 SASL plug-in

Note: Read the Code example disclaimer for important legal information.

```
import java.net.*;
import java.io.*;
import java.security.*;
import com.ibm.util.*;
import com.ibm.asn1.ASN1Exception;
import com.ibm.ldap.*;

/*
    Challenge-Response Authentication Mechanism / MD5 hash.
    See RFC 2195 ("IMAP/POP AUTHorize Extension for Simple
    Challenge/Response") and draft-ietf-ldapext-authmeth-02
    ("Authentication Methods for LDAP") for details.
*/
public class CramMD5 extends LDAPSaslBind
{
    public boolean bind(String dn, String credentials)
        throws IOException, ASN1Exception, LDAPException
    {
        String clientCreds;

        // Send initial bind request
        if (SendBindRequest("CRAM-MD5", dn, null) == true)
            return false;

        // Generate md5 hash from client's secret and server's
        // challenge and send to server
        try {
            clientCreds = "dn: " + new String(stringUTF(dn)) + " " +
                HMAC_MD5(credentials, getServerCredentials());
        } catch (NoSuchAlgorithmException nsae) {
            throw new IOException(nsae.toString());
        }
        putCredentials(clientCreds);
        return SendBindRequest();
    }
}

/*
    Hashed Message Authentication Code. See RFC 2104
    ("HMAC: Keyed-Hashing for Message Authentication")
    for details.
*/
public static String HMAC_MD5(String secret, String text)
    throws NoSuchAlgorithmException
{
    MessageDigest md5;
    byte[] ipad, opad, key;
    int i;

    // Initialize
    md5 = MessageDigest.getInstance("MD5");
    ipad = new byte[64];
    opad = new byte[64];
    key = secret.getBytes();

    // If key is larger than block size then hash key
    if (key.length > 64)
        key = md5.digest(key);

    // Perform XOR of ipad and opad with key (padded to 64 bytes).
    for (i = 0; i < key.length; ++i) {
        ipad[i] = (byte)(0x36 ^ key[i]);
        opad[i] = (byte)(0x5c ^ key[i]);
    }
}
```

```

    }
    while (i < 64) {
        ipad[i] = 0x36;
        opad[i++] = 0x5c;
    }

    // Hash ipad XOR result and text
    md5.update(ipad);
    key = md5.digest(text.getBytes());

    // Hash opad XOR result and previous hash
    md5.update(opad);
    key = md5.digest(key);

    // Return hex representation of hash (32 bytes)
    return Hex.toString(key, false);
}
}

```



Client-side caching

» Caching provides a way to store recently requested information locally. This improves performance by retrieving repeated queries locally instead of going back to a remote server for information already obtained. The down side is that there is no way to determine whether a value in cache has been changed since its retrieval. The internet draft, A Simple Caching Scheme for Lightweight Directory Access Protocol (LDAP) and X.500 Directories, addresses this issue by defining a "time-to-live" attribute that indicates the amount of time an entry can reasonably remain in cache before growing stale. The implementation of client-side caching in this class library fully supports this draft.

The algorithm for caching used in this class library is completely based on a search request. Data is not replicated locally; thus, it is not possible to query an entry two different ways from cache. For example, a search for "cn=Joe*" followed by a query for "cn=Joe Smith" does not retrieve the value from cache, although the results may reside locally as part of the "cn=Joe*" result set. Replication of LDAP data has many subtleties and is beyond the scope of this caching scheme, especially in Java^(TM). It may take more time in Java to search complicated queries locally than it would to retrieve the results from an optimized server.

Caching should not be used for data which is likely to change quickly, or when critical information is being retrieved. Caching is disabled by default in this class library, and when enabled, provides a way to bypass cache and retrieve results directly from the server. However, LDAP data is typically static and often information-based, so caching with accurate time-to-live values is a reasonable model for many applications.

When enabled, caching is used on operations that retrieve data from the server and includes the following methods:

- getAttributes
- search
- lookup
- lookupLink
- list
- listBindings
- getSchema
- getSchemaClassDefinition

Cache is created when an object of type LDAPCache is instantiated. Multiple contexts may share a single cache, meaning that a value placed in cache by one context may be retrieved from cache by another.

Caching in the IBM JNDI LDAP provider is also thread-safe. Another feature of LDAPCache is that it can be serialized to disk and later restored. The size of cache, based on the number of entries retrieved, may be changed.

A context does not use LDAPCache directly, but it relies on LDAPCacheControl to manage cache. LDAPCacheControl provides an individualized view into a shared cache. LDAPCacheControl makes it possible to adjust settings on a per request basis such as the time-to-live value, and whether to bypass cache and retrieve results directly from the server. Although LDAPCacheControl may be used on a per context basis, it can also be shared among multiple contexts. Any LDAPCacheControl reference in a context is forwarded to child contexts created by the lookup, lookupLink, and listBindings methods.

The java.naming.control.cache property is defined by the provider as a way to associate a cache control with a context. This can be done either at context creation or later through the addToEnvironment method.

To enable caching, an object of type LDAPCacheControl must be instantiated and associated with a context. The LDAPCacheControl constructor requires an LDAPCache object as a argument. The following is an example of creating a context with caching enabled:

Example: Create a context with caching enabled

Note: Read the Code example disclaimer for important legal information.

```
import com.ibm.jndi.LDAPCtx;
import com.ibm.ldap.LDAPCache;
import com.ibm.ldap.LDAPCacheControl;

LDAPCache cache = new LDAPCache();
LDAPCacheControl cacheControl = new LDAPCacheControl(cache);
env.put("java.naming.control.cache", cacheControl);
DirContext ctx = new InitialDirContext(env);
```

It is possible to adjust settings on a request basis. The settings remain in effect until they are cleared. The following forces a time-to-live value of 60 seconds:

```
cacheControl.putTTL(60);
cacheControl.putHonorServerTTL(false);
ctx.search(...);
```

The following bypasses any values stored in cache and retrieves the results directly from the server:

```
cacheControl.putReadFlag(false);
ctx.search(...);
```


Note: Because of the java.naming.batchsize property, items are added to cache until the results are completely enumerated. <<

Retrieve the IBMJNDI class version

>> The version of the Lightweight Directory Access Protocol (LDAP) classes can be obtained using the following static method:

```
String version = com.ibm.jndi.LDAPCtx.getVersion();
```

Conformance considerations and additional properties

The following known differences exist between the IBM Java[™] Naming and Directory Interface (JNDI) LDAP provider and Sun's JNDI Implementor Guidelines for LDAP Service Providers (Draft 0.2) .

1. Naming federation is not supported.
2. The use of the Secure Socket Layer (SSL) protocol is not supported by the IBM JNDI LDAP provider on iSeries servers.
3. The IBM LDAP provider should not be used to store serialized Java objects in an LDAP directory.

4. Passing URL components beyond what is needed does not result in either a `ConfigurationException` or an `InvalidNameException`.
5. The IBM JNDI LDAP provider has its own Simple Authentication and Security Layer (SASL) plug-in support. It does not support the Java SASL Application Programming Interface (API) because the API currently exists as a preview only package.
6. The provider supports unsolicited notification events, but not namespace or object change events.
7. The default value for `java.naming.referral` is `follow` and not `ignore`. When this property is set to `ignore`, the provider neither automatically adds the `ManagedSAIt` control to requests, nor does it throw a `PartialResultException` when a referral is received.

The following property is also supported:

com.ibm.jndi.ldap.so_timeout

This property defines the number of milliseconds a context block waits for data from a server. The default timeout is 5 minutes. A timeout of zero is interpreted as an infinite timeout.



The following properties have been deprecated:

- `java.naming.ldap.noBind`
- `java.naming.control.server`



JSSL

» The Java Secure Socket Layer (JSSL) is a set of Java packages that enable secure Internet communications. It implements a Java version of SSL and Transport Layer Security (TLS) protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication. Using JSSL, you can develop applications that provide for the secure passage of data between a client and a server running any application protocol (such as HTTP, Telnet, NNTP, and FTP) over TCP/IP.

For more information about JSSL, see the Java Secure Socket Extension (JSSE)  by Sun Microsystems, Inc. 

JavaMail

» The JavaMail^(TM) API provides a set of abstract classes that models an electronic (e-mail) system. The API provides general mail functions for reading and sending mail, and requires service providers to implement the protocols.

Service providers implement specific protocols. For example, Simple Mail Transfer Protocol (SMTP) is a transport protocol for sending e-mail. Post Office Protocol 3 (POP3) is the standard protocol for receiving e-mail. Internet Message Access Protocol (IMAP) is an alternative protocol to POP3.

In addition to service providers, JavaMail requires the JavaBeans Activation Framework (JAF) to handle mail content that is not plain text. This includes Multipurpose Internet Mail Extensions (MIME), Uniform Resource Locator (URL) pages, and file attachments.


All the JavaMail components are shipped as part of the IBM Developer Kit for Java. These components include the following:

- **mail.jar**


This JAR file contains JavaMail APIs, the SMTP service provider, the POP3 service provider, and the IMAP service provider.

- **activation.jar**

This JAR file contains the JavaBeans Activation Framework.

Refer to Sun Microsystems, Inc. JavaMail documentation for more information.  



Java Print Service

 The Java^(TM) Print Service API allows printing on all Java platforms. Java 1.4 provides a framework in which Java runtime environments and third parties can provide stream generator plugins for producing various formats for printing, such as PDF, Postscript, and Advanced Function Presentation (AFP). These plugins generate the output formats from bi-dimensional (2D) graphic calls.

Refer to Sun Microsystems Java Print Service documentation for more information  . 

Chapter 8. Debug programs using the IBM Developer Kit for Java

If you need to debug your Java^(TM) programs, select one of these options:

- Debug a Java program
- Debug Java and native method programs
- Debug a Java program from another display
-  Debug Java classes loaded through a custom class loader 
- Debug servlets


When you debug a Java program, your Java program is actually running in the Java virtual machine in a batch immediate (BCI) job. Your source code displays in the interactive display, but the Java program is not running there. It is running in the other job, which is a serviced job. When your Java program ends, the serviced job ends, and a message displays, stating that Job being serviced ended.

It is not possible to debug Java programs running with the Just-In-Time (JIT) compiler. If a file does not have an associated Java program, the default is to run the JIT. This can be disabled in several ways to allow debugging:

- Specify the property `java.compiler=NONE` when starting the Java virtual machine.
- Specify `OPTION(*DEBUG)` on the Run Java (RUNJVA) command.
- Specify `INTERPRET(*YES)` on the Run Java (RUNJVA) command.
- Use `CRTJVAPGM OPTIMIZATION(10)` to create an associated Java program before the Java virtual machine is started.

Note: None of these solutions affect a running Java virtual machine. If a Java virtual machine was not started with one of these alternatives, it must be stopped and restarted to be debugged.

The interface between the two jobs is established when you specify the `*DEBUG` option on the Run Java (RUNJVA) command.

For more information about the system debugger, see the WebSphere Development Studio: ILE C/C++ Programmer's Guide, SC09-2712 book  and online help information.

Debug a Java program

There are many different ways to debug a Java^(TM) program. You can use the `*DEBUG` option to view the source code before running the program. Then, you can set breakpoints, or step over or into a program to analyze errors while the program is running.

To debug Java programs, follow these steps:

1. Compile the Java program by using the `DEBUG` option, which is the `-g` option on the `javac` tool. See Debug Java programs by using the `*DEBUG` option for more details.
2. Insert the class file (`.class`) and source file (`.java`) in the same directory on your iSeries server.
3. Run the Java program by using the Run Java (RUNJVA) command on the iSeries command line. Specify `OPTION(*DEBUG)` on the Run Java (RUNJVA) command.
Note: Only a class may be debugged. If a JAR file name is entered for the `CLASS` keyword, `OPTION(*DEBUG)` is not supported.
4. The Java program source displays.

- The Work with Module List display is shown. You can add other classes and programs to debug by entering option 1 (Add program). Display their source with option 5 (Display module source).

```

-----+-----
                                Work with Module List
                                System:  AS400

Type options, press enter.
  1=Add program   4=Remove program   5=Display module source
  8=Work with module breakpoints

Opt   Program/module   Library   Type
      HELLOD           *LIBL    *SRVPGM
                        *CLASS    Selected

Command
====>
F3=Exit  F4=Prompt  F5=Refresh  F9=Retrieve  F12=Cancel
F22=Display class file name

                                Bottom
-----+-----

```

- When adding a class to debug, you may need to enter a package-qualified class name that is longer than the Program/module input field. To enter a longer name, follow these steps:
 - Enter Option 1 (Add program).
 - Leave the Program/module field blank.
 - Leave the library field as *LIBL.
 - Enter *CLASS for Type.
 - Press Enter.
 - A pop up window is displayed where you have more room to enter the package-qualified class file name.

Set breakpoints

The running of a program can be controlled with breakpoints. Breakpoints stop a running program at a specific statement.

To set breakpoints, perform the following steps:

- Place the cursor on the line of code where you would like to set a breakpoint.
- Press F6 (Add/Clear breakpoint) to set the breakpoint.
- Press F12 (Resume) to run the program.

Note: Just before the line of code runs, where the breakpoint is set, the program source is displayed indicating that the breakpoint was hit.

```

-----+-----
                                Display Module Source

Current thread:  00000019   Stopped thread:  00000019
Class file name:  HelloD
35 public static void main(String[] args)
36     {
37     int i,j,h,B[],D[] [];
38     HelloD A=new HelloD();
-----+-----

```

```

39     A.myHelloD = A;
40     HelloD C[];
41     C = new HelloD[5];
42     for (int counter=0; counter<2; counter++) {
43         C[counter] = new HelloD();
44         C[counter].myHelloD = C[counter];
45     }
46     C[2] = A;
47     C[0].myString = null;
48     C[0].myHelloD = null;

49     A.method1();
Debug . . .

F3=End program  F6=Add/Clear breakpoint  F10=Step  F11=Display variable
F12=Resume  F17=Watch variable  F18=Work with watch  F24=More key
Breakpoint added to line 41.

```

When you hit a breakpoint, if you want to set breakpoints that are only hit within the current thread, use the TBREAK command.

For more information about system debugger commands, see the WebSphere Development Studio: ILE

C/C++ Programmer's Guide, SC09-2712 book  and online help information.

For information about evaluating variables when a program stops running at a breakpoint, see Evaluate variables in Java^(TM) programs.

Step through Java programs to debug

You can step through your program while debugging. You can either step over or step into other functions. Java^(TM) programs and native methods can use the step function.

When the program source first displays, you can start stepping. The program stops before running the first statement. Press F10 (Step). Continue to press F10 (Step) to step through the program. Press F22 (Step into) to step into any function that your program calls. You can also start stepping anytime a breakpoint is hit. For information about setting breakpoints, see Set breakpoints.

```

-----+-----
                                Display Module Source
Current thread:  00000019      Stopped thread:  00000019
Class file name:  HelloD
35  public static void main(String[] args)
36  {
37      int i,j,h,B[],D[] [];
38      HelloD A=new HelloD();
39      A.myHelloD = A;
40      HelloD C[];
41      C = new HelloD[5];
42      for (int counter=0; counter<2; counter++) {
43          C[counter] = new HelloD();
44          C[counter].myHelloD = C[counter];
45      }
46      C[2] = A;
47      C[0].myString = null;
48      C[0].myHelloD = null;
49      A.method1();
Debug . . .

F3=End program  F6=Add/Clear breakpoint  F10=Step  F11=Display variable

```


1. Press F14 (Work with module list) when your Java program source displays to show the Work with Module List (WRKMODLST) display.
2. Select option 1 (Add program) to add your service program.
3. Select option 5 (Display module source) to display the *MODULE that you want to debug and the source.
4. Press F6 (Add/Clear breakpoint) to set breakpoints in the service program. For more information about setting breakpoints, see Set breakpoints.
5. Press F12 (Resume) to run the program.

Note: When the breakpoint is hit in your service program, the program stops running, and the source for the service program displays.

Debug a Java program from another display

While debugging a Java^(TM) program, the program source displays whenever it encounters a breakpoint. This may interfere with the display output of the Java program. To avoid this, debug the Java program from another display. The output from the Java program displays where the Java command is running and the program source shows on the other display.

It is also possible to debug an already running Java program in this manner as long as it is not using the Just-In-Time (JIT) compiler.

To debug Java from another display, do the following:

1. The Java program must be held, while you start setting up to debug. You can hold the Java program by making the program:
 - Wait for input from the keyboard.
 - Wait for a time interval.
 - Loop to test a variable, which requires that you set a value to eventually get the Java program out of the loop.
2. Once the Java program is held, go to another display to perform these steps:
 - a. Enter the Work with Active Jobs (WRKACTJOB) command on the command line.
 - b. Find the batch immediate (BCI) job where your Java program is running. Look under the Subsystem/Job listing for QJVACMDSRV. Look under the User listing for your User ID. Look under Type for BCI.
 - c. Enter option 5 to work with that job.
 - d. At the top of the Work with Job display, the Number, User, and Job are displayed. Enter STRSRVJOB Number/User/Job.
 - e. Enter STRDBG CLASS(classname). Classname is the name of the Java class that you want to debug. It can either be the class name that you specified on the Java command, or it can be another class.
 - f. The source for that class appears in the Display Module Source display.
 - g. Set breakpoints, by pressing F6 (Add/Clear breakpoint), whenever you would like to stop in that Java class. Press F14 to add other classes, programs, or service programs to debug. For more information about setting breakpoints, see Set breakpoints.
 - h. Press F12 (Resume) to continue running the program.
3. Stop holding your original Java program. When the breakpoints are hit, the Display Module Source display appears on the display where the Start Service Job (STRSRVJOB) command and the Start Debug (STRDBG) command were entered. When the Java program ends, a Job being serviced ended message appears.
4. Enter the End Debug (ENDDBG) command.
5. Enter the End Service Job (ENDSRVJOB) command.

Note: Ensure that you disable the Just-In-Time (JIT) when starting the Java virtual machine in the original job. This can be done with the `java.compiler=NONE` property. If the JIT runs while debugging, unexpected results may occur.

See `QIBM_CHILD_JOB_SNDINQMSG` environment variable for more information about this variable that controls whether the BCI job waits before calling the Java virtual machine.

QIBM_CHILD_JOB_SNDINQMSG environment variable

The `QIBM_CHILD_JOB_SNDINQMSG` environment variable is the variable that controls whether the batch immediate (BCI) job, where the Java^(TM) virtual machine runs, waits before starting the Java virtual machine.

If you set the environment variable to 1 when the Run Java (RUNJVA) command runs, a message is sent to the user's message queue. The message is sent before the Java virtual machine starts in the BCI job. The message looks like this:

```
Spawned (child) process 023173/JOB/QJVACMSRV is stopped (G C)
```

To view this message, enter `SYSREQ` and select option 4.

The BCI job waits until you enter a reply to this message. A reply of (G) starts the Java virtual machine.

You can set breakpoints in a `*SRVPGM` or `*PGM`, which the BCI job calls, before replying to the message.

Note: You cannot set breakpoints in a Java class, because at this point, the Java virtual machine has not been started.

Debug Java classes loaded through a custom class loader

» To debug a class loaded through a custom class loader, perform the following steps.

1. Set the `DEBUGSOURCEPATH` environment variable to the directory containing the source code, or in the case of a package-qualified class, the starting directory of the package names.

For example, if the custom class loader loads classes located under the directory `/MYDIR`, perform the following:

```
ADDENVVAR ENVVAR(DEBUGSOURCEPATH) VALUE('/MYDIR')
```

2. Add the class to the debug view from the Display Module Source screen.

If the class has already been loaded into the Java^(TM) virtual machine (JVM), just add the `*CLASS` as usual and display the source code to debug.

For example, to view the source for `pkg1/test14.class`, enter the following:

Opt	Program/module	Library	Type
1	pkg1.test14_	*LIBL	*CLASS

If the class has not been loaded into the JVM, perform the same steps to add the `*CLASS` as previously indicated. The **Java class file not available** message then displays. At this point, you may resume program processing. The JVM automatically stops when any method of the class matching the given name is entered. The source code for the class is displayed and can be debugged. <<

Debug servlets

» Debugging servlets is a special case of debugging classes loaded through a custom class loader. Servlets run in the Java^(TM) runtime of the IBM HTTP Server. One way to debug servlets is by following the instructions for classes loaded through a custom class loader.

Another way to debug a servlet is the following: <<

1. Use the `javac -g` command in the Qshell Interpreter to compile your servlet.
2. Copy the source code (.java file) and compiled code (.class file) to `/QIBM/ProdData/Java400`.
3. Run the Create Java Program (CRTJVAPGM) command against the .class file using optimization level 10, OPTIMIZE(10).
4. Start the server.
5. Run the Start Service Job (STRSRVJOB) command on the job where the servlet runs.
6. Enter `STRDBG CLASS(myServlet)`, where `myServlet` is the name of your servlet. The source should be displayed.
7. Set a breakpoint in the servlet and press F12.
8. Run your servlet. When the servlet hits the breakpoint, you can continue debugging.

Java Platform Debugger Architecture

➤ The Java^(TM) Platform Debugger Architecture (JPDA) consists of three parts:

- Java Virtual Machine Debug Interface (JVMDI)
- Java Debug Wire Protocol (JDWP)
- Java Debug Interface (JDI)

All three parts of the JPDA enable any front end of a debugger that uses the JDWP to perform debugging operations. The debugger front end can either run remotely or run as an iSeries application.

Java Virtual Machine Debug Interface

In Java^(TM) 2 SDK (J2SDK), Standard Edition, version 1.2 or higher, the Java Virtual Machine Debug Interface (JVMDI) is part of Sun Microsystems, Inc. platform application program interfaces (APIs). JVMDI allows anyone to write a Java debugger for an iSeries server in iSeries C code. The debugger does not need to know the internal structure of the Java virtual machine since it uses JVMDI interfaces. JVMDI is the lowest-level interface in JPDA that is closest to the Java virtual machine.

The debugger runs in the same multi-thread capable job as the Java virtual machine. The debugger uses Java Native Interface (JNI) Invocation APIs to create a Java virtual machine. It then places a hook at the beginning of a user class main method and calls the main method. When the main method begins, the hook is hit and debugging begins. Typical debug facilities are available, such as setting breakpoints, stepping, displaying variables, and changing variables.

The debugger handles communication between the job where the Java virtual machine is running and a job handling the user interface. This user interface is either on your iSeries server or another system.

A service program, called QJVAJVMDI that resides in the QSYS library, supports the JVMDI functions.

Java Debug Wire Protocol


The Java Debug Wire Protocol (JDWP) is a defined communication protocol between a debugger process and the JVMDI. JDWP can be used from either a remote system or over a local socket. It is one layer removed from the JVMDI, but is a more complex interface.

Start JDWP in QShell

To start JDWP and run the Java class `SomeClass`, enter the following command in QShell:

```
java -interpret -Xrunjdpw:transport=dt_socket,  
address=8000,server=y,suspend=n SomeClass
```

In this example, JDWP listens for connections from remote debuggers on TCP/IP port 8000, but you can use any port number you want; `dt_socket` is the name of the SRVPGM that handles the JDWP transport and does not change.

For additional options that you can use with `-Xrunjdpw`, see [Sun VM Invocation Options](#)  by Sun Microsystems, Inc.

Start JDWP from a CL command line

To use the `-Xrun` option with the CL command, the `os400.xrun.option` property can be defined to be the same string that you would have used on the QShell command line. To start JDWP and run the Java class `SomeClass`, enter the following command:



```
JAVA CLASS(SomeClass) INTERPRET(*YES)
PROP((os400.xrun.option 'jdpw:transport=dt_socket,address=8000,
server=y,suspend=n'))
```

Many JVMDI functions do not work at optimization levels 10 and 20. Therefore, it is recommended that your applications run with the interpreter since all functions work with it.

Java Debug Interface


Java Debug Interface (JDI) is a high-level Java language interface provided for tool development. JDI hides the complexity of JVMDI and JDWP behind some Java class definitions. JDI is included in the `rt.jar` file, so the front end of the debugger exists on any platform that has Java installed.

If you want to write debuggers for Java, you should use JDI because it is the simplest interface and your code is platform-independent.

For more information on JDBA, see [Java Platform Debugger Architecture Overview](#)  by Sun Microsystems, Inc. 

Find memory leaks

» ANZJVM finds object leaks by taking two copies of the garbage collection heap that are separated by a specified time interval. To find object leaks, you would look at the number of instances of each class in the heap. Classes that have an unusually high number of instances should be noted as possibly leaking.

You should also note the change in number of instances of each class between the two copies of the garbage collection heap. If the number of instances of a class continually increases, that class should be noted as possibly leaking. The longer the time interval between the two copies, the more certainty you have that objects are actually leaking. By running ANZJVM a series of times with a larger time interval, you should be able to diagnose with a high degree of certainty what is leaking. 

Chapter 9. Troubleshooting the IBM Developer Kit for Java

If you encounter problems while using the IBM Developer Kit for Java^(TM), take any of these steps to determine the source of the problem.

- **»** You may notice some limitations when using the IBM Developer Kit for Java. This topic identifies any known limitations, restrictions, or unique behaviors. **«**
- Find the job log from the job that ran the Java command. Also, look for the job log from the batch immediate (BCI) job where the Java program ran to analyze the cause of failure.
- Collect useful data for an authorized program analysis report (APAR).
- Apply program-temporary fixes (PTFs).
- Know how to get support if you detect a potential defect in the IBM Developer Kit for Java.

Limitations

When using the IBM Developer Kit for Java^(TM), you may notice that there are some limitations in how it is used. This list identifies any known limitations, restrictions, or unique behaviors.

- When a class is loaded and its superclasses are not found, the error indicates that the original class was not found. For example, if class B extends class A, and class A is not found when loading class B, the error indicates that class B was not found, even though it is actually class A that was not found. When you see an error that indicates that a class was not found, check to make sure that the class and all of its superclasses are in the CLASSPATH. This also applies to interfaces that are implemented by the class being loaded.
- **»** The garbage collection heap is limited to 132 GB.
- The number of constructed objects is limited to 132 million. **«**
- The java.net backlog parameter on an iSeries server may behave differently than on other platforms. For example:
 - Listen backlogs 0, 1
 - Listen(0) means to allow one pending connection; it does not disable a socket.
 - Listen(1) means to allow one pending connection, and means the same as Listen(0).
 - Listen backlogs > 1
 - This allows many pending requests to remain on the listen queue. If a new connection request arrives and the queue is at the limit, then it deletes one of the pending requests.
- **»** You can only use the Java virtual machine, regardless of the JDK version you are using, in multi-thread capable (that is, thread-safe) environments. The iSeries server is thread-safe, but some file systems are not. The integrated file system topic has a list of nonthread-safe file systems. **«**

Find job logs for Java problem analysis

Use the job log from the job that ran the Java^(TM) command, and the batch immediate (BCI) job log where the Java program ran, to analyze causes of a Java failure. They both may contain important error information.

There are two ways to find the job log for the BCI job. You can find the name of the BCI job that is logged in the job log of the job that ran the Java command. Then, use that job name to find the job log for the BCI job.

You can also find the job log for the BCI job by following these steps:

1. Enter the Work with Submitted Jobs (WRKSBMJOB) command on the iSeries command line.
2. Go to the bottom of the list.

3. Look for the last job in the list, called QJVACMDSRV.
4. Enter option 8 (Work with Spooled Files) for that job.
5. A file called QPJOBLOG displays.
6. Press F11 to see view 2 of the spooled files.
7. Verify that the date and time match the date and time when the failure occurred.
Note: If the date and time do not match the date and time when you signed off, continue looking through the list of submitted jobs. Try to find a QJVACMDSRV job log with a date and time that matches when you signed off.

If you are unable to find a job log for the BCI job, one may not have been produced. This happens if you set the ENDSEP value for the QDFTJOB job description too high or the LOG value for the QDFTJOB job description specifies *NOLIST. Check these values, and change them so that a job log is produced for the BCI job.

To produce a job log for the job that ran the Run Java (RUNJVA) command, perform the following steps:

1. Enter SIGNOFF *LIST.
2. Then, sign back on.
3. Enter the Work with Spooled Files (WRKSPLF) command on the iSeries command line.
4. Go to the bottom of the list.
5. Find a file named QPJOBLOG.
6. Press F11.
7. Verify that the date and time match the date and time when you entered the signoff command.
Note: If the date and time do not match the date and time when you signed off, continue looking through the list of submitted jobs. Try to find a QJVACMDSRV job log with a date and time that matches when you signed off.

Collect data for Java problem analysis

To collect data for an authorized program analysis report (APAR), follow these steps:

1. Include a complete description of the problem.
2. Save the Java[™] class file that caused the problem while running.
3. You can use the SAV command to save objects from the integrated file system. You may need to save other class files that this program must run. You may also want to save and send in an entire directory for IBM to use when trying to reproduce the problem, if necessary. This is an example of how to save an entire directory.

Example: Save a directory

Note: Read the Code example disclaimer for important legal information.


```
SAV DEV('/QSYS.LIB/TAP01.DEVD') OBJ(('mydir'))
```

If possible, save the source files for any Java classes that are involved in the problem. This is helpful to IBM when reproducing and analyzing the problem.

4. Save any service programs that contain native methods that are required to run the program.
5. Save any data files that are required to run the Java program.
6. Add a complete description of how to reproduce the problem. This should include:
 - The value of the CLASSPATH environment variable.
 - A description of the Java command that was run.
 - A description of how to respond to any input that is required by the program.
7. Include any vertical licensed internal code (VLIC) logs that have occurred near the time of failure.

8. Add the job log from both the interactive job and the BCI job where the Java virtual machine was running.

Get support for the IBM Developer Kit for Java


Support services for the IBM Developer Kit for Java^(™) are provided under the usual terms and conditions for iSeries software products. Support services include program services, voice support, and consulting services. Use the online information that is provided at IBM iSeries Home Page  under the topic "Support" for more information. Use IBM Support Services for 5722-JV1 (IBM Developer Kit for Java). Or, contact your local IBM representative.

You may, at IBM direction, be required to obtain a more current level of the IBM Developer Kit for Java to receive Continued Program Services. For more information, see Support for multiple Java Development Kits (JDKs).

Resolving defects of the IBM Developer Kit for Java program are supported under program services or voice support. Resolving application programming or debugging issues are supported under consulting services.

The IBM Developer Kit for Java application program interface (API) calls are supported under consulting services, unless:

1. It is clearly a Java API defect as demonstrated by re-creation in a relatively simple program.
2. It is a question that asks for documentation clarification,
3. It is a question about the location of samples or documentation.

All programming assistance is supported under consulting services. This includes the program samples that are provided in the IBM Developer Kit for Java licensed program (LP) product. Additional samples may be available on the Internet at IBM iSeries Home Page  on an unsupported basis.

The IBM Developer Kit for Java LP provides information about solving problems. If you believe that there is a potential defect in the IBM Developer Kit for Java API, a simple program that demonstrates the error is required.

Chapter 10. Code examples for the IBM Developer Kit for Java

The following is a list of code examples for the IBM Developer Kit for Java^(TM).

Note: Read the Code example disclaimer for important legal information.

CL commands

- >> ANZJVM<<
- CHGJVAPGM
- CRTJVAPGM
- DLTJVAPGM
- DMPJVM
- DSPJVAPGM
- RUNJVA

Internationalization

- DateFormat
- NumberFormat
- ResourceBundle

JDBC

- >> Access property
- Blob<<
- CallableStatement interface
- >> Change values with a statement through another statement's cursor
- Clob
- Create a UDBDataSource and bind it with JNDI
- Create a UDBDataSource, and obtain a user ID and password
- Create a UDBDataSourceBind and set DataSource properties<<
- DatabaseMetaData interface
- >> Create a UDBDataSource and bind it with JNDI
- Datalink
- Distinct types<<
- Embed SQL Statements
- >> End a transaction
- Invalid user ID and password
- JDBC
- Multiple connections that work on a transaction
- Obtain an initial context before binding UDBDataSource
- ParameterMetaData
- Remove values from a table through another statement's cursor<<
- ResultSet interface
- >> ResultSet sensitivity
- Sensitive and insensitive ResultSets

- Set up connection pooling with UDBDataSource and UDBConnectionPoolDataSource
- SQLException<<
- >> Suspend and resume a transaction
- Suspended ResultSets
- Test the performance of connection pooling
- Test the performance of two DataSources
- Update BLOBs
- Update CLOBs
- Use a connection with multiple transactions
- Use BLOBs
- Use CLOBs
- Use DB2CachedRowSet properties and DataSources
- Use DB2CachedRowSet properties and JDBC URLs
- Use JTA to handle a transaction
- Use metadata ResultSets that have more than one column
- Use native JDBC and Toolbox JDBC concurrently
- Use PreparedStatement to obtain a ResultSet
- Use the execute(Connection) method to use an existing database connection
- Use the execute(int) method to batch database requests together
- Use the populate method
- Use the setConnection(Connection) method to use an existing database connection
- Use the Statement object's executeUpdate method<<

Java Authentication and Authorization Service

- JAAS HelloWorld example
- JAAS SampleThreadSubjectLogin example

>> Java Generic Security Service

- Sample non-JAAS client program
- Sample non-JAAS server program
- Sample JAAS-enabled client program
- Sample JAAS-enabled server program <<

>> Java Naming and Directory Interface

- Add an entry in a directory
- Delete an entry in a directory
- Add an entry in a directory
- Rename a directory entry
- Specify a list of binary attribute names<<

>> Java Secure Sockets Extension

- SSL client and server using an SSLContext object<<

>> Java with other programming languages

- Call a CL program

- Call a CL command
- Call another Java program
- Call Java from C
- Call Java from RPG
- Input and output streams
- Invocation API
- OS/400 PASE native method for Java <<
- Sockets

>> Optional packages

- JCE<<

Performance tools

- Java Performance Data Converter

>> Run a host without a GUI

- Set up Remote AWT<<

SQLJ

- Embed SQL Statements in your Java application

Secure sockets layer

- Socket factories
- Server socket factories
- Secure sockets layer
- Secure sockets layer server

Chapter 11. IBM Developer Kit for Java reference

➤ The following are references for the Developer Kit for Java^(TM):

Javadoc

- iSeries-specific JAAS Javadoc
- JAAS API Specification

Java 2 Platform, Standard Edition, version 1.3.1

- Java 2 Platform, Standard Edition, v1.3.2 API Specification
- Abstract Window Toolkit (AWT)
- Java IDL
- Input Method Framework
- Internationalization
- JDBC API
- JNI - Java Native Interface
- Java Remote Method Invocation (RMI)
- RMI - Remote Method Invocation
- Security
- Java 2 SDK Tools<<

Code disclaimer information

This document contains programming examples.

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

All sample code is provided by IBM for illustrative purposes only. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

All programs contained herein are provided to you "AS IS" without any warranties of any kind. The implied warranties of non-infringement, merchantability and fitness for a particular purpose are expressly disclaimed.



Printed in U.S.A.