



INTRODUCTION TO
SQLJ

- Revised January 2007 -

1 What is SQLJ?

SQL for Java (SQLJ) is a way to seamlessly embed SQL syntax directly in the Java programming language. It uses several classes included in the JDBC package but also have some classes of its own. This is done through an additional pre-compiler that interprets the SQLJ specific syntax and creates source files in Java. These files are then compiled as any other Java application through the *javac* compiler.

To be able to compile and execute a Java program that uses the SQLJ framework you need to have three things. First, you need to have the *J2SE Development Kit (JDK)* installed. This environment is free of charge and could be downloaded from: <http://java.sun.com/j2se/1.5.0/download.jsp>. Secondly, you need to have the SQLJ pre-compiler (*sqlj.exe*) and the SQLJ Java package (*sqlj.zip*, usually found under *C:\Program Files\IBM\SQLLIB\java* or where you have installed DB2). Last, you need to have the program *nmake.exe* and the batch-file *embprep.bat* (placed in *c:\MyProg* at your prepared environment)¹.

When writing a SQLJ application you save the source code as a *.sqlj* file. To produce a Java source file out of this you use the pre-compiler. The resulting *.java* file can then be compiled through the default Java compiler *javac*. This produces the *.class* file that could be executed by the virtual machine. The *embprep.bat* file is especially made for your course and need to be modified in case it is used for other purposes.

2 Makefile

A makefile is a set of instructions on how to compile an application or perform a certain sequence of commands. In these instructions you are able to define a sequence of commands that are to be executed when triggering the *nmake* application. Consider the following:

```
crazy:
  cls
  dir
```

Note! Every command must be entered on a single row and started with a tab character.

When executing the command *nmake crazy* in a Windows Command Prompt it'll issue the command *cls* followed by *dir*. It is also possible to declare parameters in a *makefile* by using the commands:

```
DIRNAME=sample
```

DIRNAME is the name of the parameter and *sample* is its value. If you want to use this parameter you type:

```
crazy:
  cls
  dir $(DIRNAME)
```

This would issue the command *cls* followed by *dir sample*

In makefiles it is also possible to use condition constructs. You can define files that have to be available before an attempt to access them is made. Do this by adding the file names according to the example below:

¹ These two files and the makefile are actually not necessary, but they are convenient. They just offer a simpler command interface than having to write several commands every time.

```
crazy : file1.txt file1.bak
cls
dir $(DIRNAMN)
```

If you use the statement *nmake crazy*, the *nmake* application will try to find *file1.txt*. If the file is missing the application will try to create it according to any available instructions. If no such instructions exist it will return a "Fatal: 'file1.txt' does not exist – don't know how to make it" error. The following example shows how to specify instructions for file creation:

```
# Create the constant DN
DN=sample

#Specifying the crazy command
crazy : file1.txt file1.bak
dir $(DN)

#Instructions to create file1.txt
file1.txt :
dir >file1.txt

#Instructions to create file1.bak
file1.bak : file1.txt
copy file1.txt file1.bak
```

This will set the constant *DN* to the value *sample* and test that the files *file1.txt* and *file1.bak* exists before trying to issue the command *dir sample*. If one of the files is missing it will be made by piping a *dir* command into *file1.txt* and by copying this file into *file1.bak*. Every row that starts with the *#* character is treated as a remark.

3 An example (SQLJTest):

On the location [\\db-srv-1\StudentCourseMaterial\IS4_spring_2007\SQLJTest](#) on the local network you'll find the file *SQLJTest.sqlj* and the *makefile*. *SQLJTest.sqlj* is the complete source code of this example and there you'll see all syntax used in this introduction in its context. *makefile* is the file containing the instructions used together with the *nmake* program to create your running Java application.

1. First of all, create a new folder (preferably on your d: drive) and name it something convenient like for example *SQLJTest*.
2. Copy the source file *SQLJTest.sqlj* and the *makefile* from the network to your new folder.

The *SQLJTest.sqlj* file is not compiled yet, as you'll see if you open it in a text editor like for example notepad. In fact it's not even a complete Java source file. If you try to compile it through *javac*, all you'll get is a handful of problems. *SQLJTest.sqlj* is still in its SQLJ specific high level programming language, which is easily readable and meaningful for the human eye.

3. Compile the *.sqlj* file by using the *SQLJ* pre-compiler to create a new *.java* file. Probably the easiest way to do this is by placing yourself in the newly created folder using a *Windows Command Prompt* (in this case *d:\SQLJTest*) and issue the command: **nmake SQLJTest**. This will execute the following commands from the *makefile*:

```
# Use the java compiler
CC= javac
```

```
# Other variables
DATASOURCE1=sample
TESTUID=
TESTPWD=

SQLJTest.java : SQLJTest.sqlj
               sqlj SQLJTest.sqlj
SQLJTest.class : SQLJTest.java SQLJTest_SJProfile0.ser
SQLJTest      : SQLJTest.class
               $(CC) SQLJTest.java
               embprep SQLJTest $(DATASOURCE1) $(TESTUID) $(TESTPWD)
```

To run the application we need to have the *SQLJTest.class* file to address the Java virtual machine. If *SQLJTest.class* is missing we need to have the *SQLJTest.java* and *SQLJTest_SJProfile0.ser* files. *SQLJTest.java* can be created if the *SQLJTest.sqlj* file is available. The three constants that are used (DATASOURCE1, TESTUID, TESTPWD), specifies the database information needed in the final application.

The source code specified in the *.sqlj* files is very similar to plain Java. All that differs are the database specific commands which in SQLJ are denoted by a beginning *#sql* markup on each row. These are the rows the pre-compiler turn into Java when turning the file from *.sqlj* to *.java*.

3.1 Explanation of the source code

Open up *SQLJTest.sqlj* to view the application source code. Exactly as the SQLJ pre-compiler, Java compiler and the virtual machine, we'll go through the file, starting at the top.

```
import sqlj.runtime.*;
import sqlj.runtime.ref.*;
...
```

These are statements that include already compiled Java methods and constructs from other packages. The star '*' states that we shall include all classes in the *sqlj.runtime* and *sqlj.runtime.ref* package paths.

```
...
#sql iterator SQLJTest_Cursor1 (String empno, String firstnme);
#sql iterator SQLJTest_Cursor2 (String);
...
```

These statements define two SQLJ data types that later can be used to define cursor variables. The first one is called *SQLJTest_Cursor1* and the other is called *SQLJTest_Cursor2*. *SQLJTest_Cursor1* defines the two string arguments as *empno* and *firstnme* while *SQLJTest_Cursor2* only defines the argument as a string value.

Further down in *SQLJTest.sqlj* is the main method that starts with defining two cursor handles that are defined just like any other Java variable with the *SQLJTest_Cursor* data types we specified in the previous step:

```
...
SQLJTest_Cursor1 cursor1;
SQLJTest_Cursor2 cursor2;
...
```

The next part handles the database connection. Note the URL *jdbc:db2:sample* specifying the path to the DB2 database *sample*. Also note the use of *DefaultContext* that specifies the

database connection and handles the execution of SQL queries, and the *con* variable that supplies a handle to this connection. If the *con* variable is set to auto commit the queries, they are applied directly into the database when carried out. If set to false, an explicit *commit* command has to be issued for the changes to be effective.

```
...
#sql cursor1 = {SELECT empno, firstnme FROM db2admin.employee};
...
```

This select statement has to fetch the same fields as the arguments specified in *SQLJTest_Cursor1*. The fetched rows have to have just as many columns, have the same column names in the exact same order as the arguments given by *SQLJTest_Cursor1*. Using any other construct will fail to retrieve the information we want to collect. Since *SQLJTest_Cursor1* is specified to take two string arguments (*empno* and *firstnme*) using *cursor1* that's specified to this type works just fine. However, using the *cursor2* variable wouldn't work. To pull out the information from the data collection after it has been fetched, we can do the following:

```
...
while (cursor1.next()) {
    str1 = cursor1.empno();
    str2 = cursor1.firstnme();
    ...
}
cursor1.close();
...
```

The WHILE-statement iterates through the whole collection of data. In the loop we retrieve the result row after row, accessing the row's column by specifying a method in the cursor called after their name. When the data collection is exhausted, the *next* method returns false and the *close* method closes the connection to the database.

Simpler SQL queries that return no or only a single row of data don't have to use cursors. Consider the following:

```
...
#sql {UPDATE db2admin.employee SET firstnme ='SHILI' WHERE empno =
'000010'};
...
```

This would update the employee number 000010 by setting his/her first name to SHILI. It's also possible to denote variables inside the query by using the colon qualifier character ':' in front of the variable name. Putting the integer value from a select statement into the variable *count1* would look like this:

```
...
#sql { SELECT count(*) into :count1 FROM db2admin.employee };
...
```

Using undeclared cursors like the *SQLJTest_Cursor2* with unspecified argument names needs another approach. In this case you could use the *FETCH* statement that iterates through a collection of data and puts the retrieved values into the given variable. In the example below the first names gathered from the select statement are put into the *str2* variable. The *endFetch* method ends the loop when there are no more rows in the data collection:

```
...
str1 = "000010";
#sql cursor2 = { SELECT firstnme from db2admin.employee where empno =
```

```

:str1 };
...
while (true) {
    #sql { FETCH :cursor2 INTO :str2 };
    if (cursor2.endFetch()) break;
    ...
}
cursor2.close();
...

```

If you like, SQLJ could also perform transaction management. To rollback a transaction, use the command:

```

...
#sql { ROLLBACK work };
...

```

Basically, using SQLJ is designed to be easy. Creating a Java application from your SQLJ file should be as easy as running the command *nmake SQLJTest* to compile it and *java SQLJTest* to execute it. This would look like this:

D:\SQLJTest>nmake SQLJTest

```

IBM(R) Program Maintenance Utility for Windows(R)
Version 3.50.000 Feb 13 1996
Copyright (C) IBM Corporation 1988-1995
Copyright (C) Microsoft Corp. 1988-1991
All rights reserved.

```

```

    sqlj SQLJTest.sqlj
    javac SQLJTest.java
    embprep SQLJTest sample

```

```

[ibm][db2][jcc][sqlj]
[ibm][db2][jcc][sqlj] Begin Customization
[ibm][db2][jcc][sqlj] Loading profile: SQLJTest_SJProfile0
[ibm][db2][jcc][sqlj] Customization complete for profile
SQLJTest_SJProfile0.ser

[ibm][db2][jcc][sqlj] Begin Bind
[ibm][db2][jcc][sqlj] Loading profile: SQLJTest_SJProfile0
[ibm][db2][jcc][sqlj] Driver defaults(user may override): BLOCKING ALL
VALIDATE
BIND STATICREADONLY YES
[ibm][db2][jcc][sqlj] Fixed driver options: DATETIME ISO DYNAMICRULES BIND
[ibm][db2][jcc][sqlj] Binding package SQLJTE01 at isolation level UR
[ibm][db2][jcc][sqlj] Binding package SQLJTE02 at isolation level CS
[ibm][db2][jcc][sqlj] Binding package SQLJTE03 at isolation level RS
[ibm][db2][jcc][sqlj] Binding package SQLJTE04 at isolation level RR
[ibm][db2][jcc][sqlj] Bind complete for SQLJTest_SJProfile0

```

D:\SQLJTest>java SQLJTest

Retrieve some data from the database...

Received results:

```

empno= 000010 firstname= CHRISTINE
empno= 000020 firstname= MICHAEL
empno= 000030 firstname= SALLY
empno= 000050 firstname= JOHN
empno= 000060 firstname= IRVING
empno= 000070 firstname= EVA
empno= 000090 firstname= EILEEN
empno= 000100 firstname= THEODORE
empno= 000110 firstname= VINCENZO

```

```
empno= 000120 firstname= SEAN
empno= 000130 firstname= DOLORES
empno= 000140 firstname= HEATHER
empno= 000150 firstname= BRUCE
empno= 000160 firstname= ELIZABETH
empno= 000170 firstname= MASATOSHI
empno= 000180 firstname= MARILYN
empno= 000190 firstname= JAMES
empno= 000200 firstname= DAVID
empno= 000210 firstname= WILLIAM
empno= 000220 firstname= JENNIFER
empno= 000230 firstname= JAMES
empno= 000240 firstname= SALVATORE
empno= 000250 firstname= DANIEL
empno= 000260 firstname= SYBIL
empno= 000270 firstname= MARIA
empno= 000280 firstname= ETHEL
empno= 000290 firstname= JOHN
empno= 000300 firstname= PHILIP
empno= 000310 firstname= MAUDE
empno= 000320 firstname= RAMLAL
empno= 000330 firstname= WING
empno= 000340 firstname= JASON
```

```
Retrieve the number of rows in employee table...
There are 32 rows in employee table.
```

```
Update the database...
Retrieve the updated data from the database...
Received results:
```

```
empno= 000010 firstname= SHILI
```

```
Rollback the update...
Rollback done.
```

4 Writing your first SQLJ application

It's greatly recommended to use your own folder under the *d:* drive as a base when creating your application. Try to reuse as much of the *SQLJTest.sqlj* and the *makefile* as possible, modify them and make adjustments rather than starting from scratch. A few points to keep in mind when doing this:

1. The name of the file and the class must be the same. Renaming the class = renaming the file and vice versa.
2. The instructions in the makefile must correspond to the filenames in use.
3. The database information in the makefile must correspond to the database used by your program.
4. The path to the database and the login information must be correct both in the makefile and in the application.
5. The arguments specified in the declared SQLJ cursors must correspond to the fields in the SQL queries.