

DB2 Tutorial

This tutorial will introduce you to our DB2 setup and the joy of database table creation using DB2. This document assumes that you are sitting on one of the CASLAB machines. Not everything will be applicable to those working at home (although we've made a few notes to help you get your home installation working).

One note before we begin. A “**schema**” in DB2 refers to a qualifying name used to group objects together within a database. (Note this is different from the schemas we have been discussing in class). All tables have a fully qualified name consisting of its schema name and the table name. If the schema we choose is, say, “Finance” and the tables we create are called Employee and Department, then the fully qualified name for each table is “*Finance.Employee*” and “*Finance.Department*”. Schemas allow us to create 2 tables with the same name in the same database. This is important for CISC 332 where many groups are using the same physical database and perhaps creating tables by the same name. For CISC 332 you will use “A” plus your group id as your “schema” name when you create tables. So, if your group id is 10, then use “A10” as your schema name.

DB2 Control Center

Let's first take a look at the DB2 GUIs. Although you won't use these much during the course of CISC 332, they will help you to understand the CISC 332 database set up.

The most useful tool is the DB2 “*Control Center*”. As the name implies, this center is used for the control of databases known to your DB2 installation. In CASLAB, only a DB2 client is installed. No local databases can be created, but the client may be used to access remote databases.

Start the Control Center up by choosing (something resembling)

IBM DB2 -> DB2 Copy 1 -> General Admin Tools->Control Center from the Start menu. It will take a while for the control center to come up on slower machines. Be patient!

On the left panel you will see a folder labeled “*All Databases*”. Expand this folder icon so that you can see the databases listed. You will likely be asked for a login & password for validation. Use your CASLAB userid & password (the same one that you used to gain access to the local machine). You will find here, among others, a list of 10 databases called “*CISC332A, CISC332B ... CISC332J*”. These databases reside on the unix server named “*cronus*” and are “catalogued” for use by the local machine, meaning that you can access this database from the local machine. These are the databases that you will use for the course project. The database that your group will use corresponds to the “tens” digit of your group id. For instance, if you are group 1, 11, 21 etc , use cisc332a, 2, 12, 22 etc, use cisc332b etc. (It is not terribly important which database you use, just choose one; the important thing is that we want the load to be spread out among the class). Note that several groups will use the same database, but you have access to

only your tables. For the purpose of this lab, if you don't know your group number, just use your initials or first name for your schema.

Home Users: If you are at home you may encounter an error message that reads "No start database manager command was issued". This means that the DB2 process is not running. From the control center, you can start the db2 processes by right-clicking on the "DB2" icon near the top of the tree structure in the left panel. You will find a "start" option. Choose this. Note that in the lab, students do not have the authority to start/stop DB2 processes. They are always running by default.

Click on one of the CISC332 database icons to expand the tree further. Here you will see a listing of all the objects associated with the database; Tables, Views etc. The ones we will be most concerned with in this course are Tables and Views. Double click the "*Tables*" icon. In the right-hand pane you will see a listing of tables associated with this database. Notice how many there are, and we haven't created any tables yet! At this point, the tables that you see are all catalog tables that contain meta-data descriptions of the database (such as how many tables are associated with the database, the attribute names of each table & the data types for each of the tables and so on).

Home Users: If you are working from home, you will need to create a database before continuing further. You can do this by right clicking on the "database" icon. Just provide a name for your database, the other parameters are unimportant at this stage.

Now return to icons seen in the left pane. Open the icon for the local system and find the database called "*Sample*". This database is a demo database that is shipped with DB2. We use this for the SQL assignment in CISC 332. This database resides on each local machine in CASLAB. Find the Employee table and explore it by double clicking on the table name. Find out what attributes are associated with it. What are the domains of the attributes? Which columns can contain null values?

Command Line DB2

In CISC 332 you will be using the command line, or scripts (batch files) that are executed in order to create tables, drop tables and insert data. First we'll examine the use of the command line to connect to a database and to create a table. DB2 has a command window (db2cmd) which is similar to an MS DOS window. Be careful not to confuse the two as DB2 commands will not run properly in an MS DOS window.

To start the DB2 command window choose [IBM DB2->DB2 Copy1->Command Line Tools->Command Window](#) from the Start menu. You should see the MS-DOS-like window appear.

To see what databases are available for your use type:

```
db2 list database directory
```

This should return a listing of databases that looks something like this (your list will be much longer):

```
System Database Directory
```

Number of entries in the directory = 13

Database entry:

Database alias	= CISC332A
Database name	= CISC332A
Database drive	=
Database release level	= 9.00
Comment	=
Directory entry type	= Remote
Catalog node number	= 0

Note that all commands that you type in the DB2 Command window begin with the keyword "db2", thus sending the command to the db2 process for interpretation. If you wish, you can enter the command line interpreter by typing db2 <return>. This provides you with a prompt at which you can type commands on the command line without prefacing them with "db2".

To connect to one of the CISC 332 databases, say, CISC332c type

`db2 connect to cisc332c user youruserid`

where youruserid is your caslab user id. You will be prompted for a password. Type in your CASLAB password. You will see something like this:

Database Connection Information

Database server	= DB2/NT 9.1.0
SQL authorization ID	= WENDY
Local database alias	= CISC332C

Being "*connected*" means that you can now type in commands that allow you to examine the objects currently associated with the database, or you can create new objects in the database.

Let's create a table using the *CREATE TABLE* command. Issue the following command to see what tables are currently associated with you (or your schema):

`db2 list tables for schema <schema_name>`

where <schema_name> corresponds to the schema that was discussed earlier (A10 if you happen to be group 10).

You should see:

Table/View	Schema	Type	Creation time

0 record(s) selected.			

If you see something like the following instead:

Table/View	Schema	Type	Creation time
Employee	A10	T	10:50

1 record(s) selected.

Then, type: `db2 drop table A10.Employee` to drop the table.

To create an Employee table, issue the following:

```
db2 create table Employee (SIN char(9) not null primary key, DOB date)
```

This command has created a table with the default schema (which is your user id). The default schema is used whenever you do not explicitly use the schema name. Issue a `db2 list tables` command to see the table you created. This command assumes your default schema, your user id, unless you qualify with "FOR SCHEMA <schema_name>".

Drop this table now using

`db2 drop table Employee`

If you wish to create a table using a specific schema (such as your first name), qualify the table name with the schema name when you create it by replacing XXX with your schema name:

```
db2 create table XXX.Employee (SIN char(9) not null primary key, DOB date)
```

Use the "`db2 list tables for schema <schema_name>`" command to view the table you just created. You have just created 2 tables called Employee. How can they both exist in the same database? Also, many of your classmates have each created 2 tables called "Employee" in the same database as your tables! Think "schema"!

Drop the table created under your default schema using:

`db2 drop table Employee`

Now insert some data into one of your tables. Here's an example:

```
db2 insert into <schema_name>.Employee values ('123456789', '12-12-1994')
```

Check out what's in your table using "`db2 select * from <schema_name>.Employee`" and you should see the data that you just inserted.

Ok, we have now created a table using the command line. Use the "`db2 drop table <schema_name>.Employee`" command to drop the table you have just created.

Using Batch Files

The simplest way to create and populate your tables is through a batch file containing many SQL statements separated by a ";". (This is one way to do it, there are many other methods that you might use). The use of batch files will save you many keystrokes as you will find that you need to re-run your table creations several times before actually getting it right. If you do all your work via the command line or a GUI, you will type & re-type the same commands over and over. With a batch file, you can make modifications and re-run the commands easily. It also provides you with a file that can be handed in for marking purposes.

A sample batch file is found at:

<http://www.cs.queensu.ca/home/cisc332/documentation.html>.

Click on the download for the *Company* database and save it to your z: drive. (You may have to use Mozilla Firefox to download this file...there were some problems with this last year). Unzip this file onto your z: drive, and you will find 3 files; *Company.ddl*, *Query.ddl* and a readme file. Modify the *Company.ddl* file changing the database name to the CISC332 database that corresponds to your group id (A10 if you happen to be group 10). While you are here, take a look at what *Company.ddl* does. Follow the instructions in the readme file to build the company database using the batch file.

You can use the company batch file as a template for your own project. Please note the *connect* and the *terminate* commands at the beginning and the end of the file. Please ensure that if you have a "connect" in your script that you also have a "terminate" as each connection uses resources and you may eventually may block others from connecting.

The query.ddl file will run some sample SQL queries against the Company database. You will find this useful for our SQL unit.

You have now completed the DB2 tutorial!