

How-To: Use Database Tools and Utilities in DB2

IBM Data Studio

IBM Data Studio is included in every DB2 edition. IBM Data Studio provides a single integrated environment for database administration and application development. You can perform tasks that are related to database modeling and design, developing database applications, administering and managing databases, tuning SQL performance, and monitoring databases all in one single tool. It is an ideal tool that can greatly benefit a team environment with different roles and responsibilities.

IBM Data Studio comes in three flavors: full client, administration client, and web console.

The full client includes both the database administrative and the application development capabilities. The development environment is Eclipse-based. This offers a collaborative development environment by integrating with other advanced Eclipse-based tools such as InfoSphere Data Architect and InfoSphere Optim pureQuery Runtime. Note that some of the advanced InfoSphere tools are only included in the DB2 Advanced editions and the DB2 Developer Edition. You can also separately purchase the advanced tools.

The administration client is a subset of the full client. It still provides a wide range of database administrative functionality such as DB2 instance management, object management, data management, and query tuning. Basic application development tasks such as SQL Builder, query formatting, visual explain, debugging, editing, and running DB2 routines are supported. Use the full client for advanced application development features.

The web console, as the name implies, it is a web-based browser interface that provides health monitoring, job management, and connection management.

NOTE

IBM Data Studio also provides collaborative database development tools for DB2 for z/OS, DB2 for i, Informix, and other non-IBM databases. To see a list of IBM Data Studio features by data server, refer to the documentation at <http://www.ibm.com/support/docview.wss?uid=swg27022148>.

IBM Data Studio Workspace and the Task Launcher

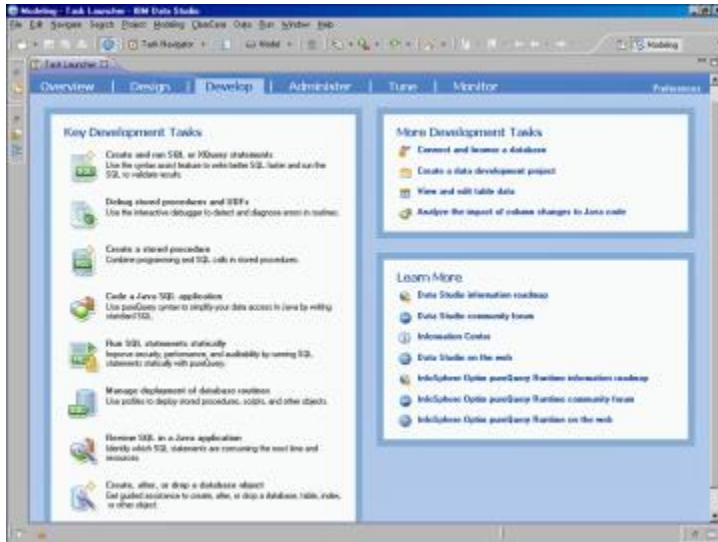
When you have successfully installed the IBM Data Studio, you are asked to provide a workspace name.

A **workspace** is a folder that saves your work and projects. It refers to the desktop development environment, which is an Eclipse-based concept.

Task Launcher is displayed, which highlights the following category of tasks:

- Design
- Develop
- Administer
- Tune
- Monitor

Each category is described in more detail in its own tab. Click any tab, and you see the key and primary tasks listed in the box on the left. See [Figure 4.26](#) to get an idea on how to navigate the Task Launcher.



[Figure 4.26](#) The IBM Data Studio Task Launcher

As an example, the figure shows you the Develop tasks. You can find the key development tasks on the left. On the top right, it lists more tasks related to development. On the bottom right, IBM Data Studio provides a few documentation links where you can learn more about development. Where appropriate, it also suggests the advanced tools available in the InfoSphere Optim portfolio that apply to the task you have selected.

Connection Profiles

Every task you were to perform against a database requires to first establish a database connection. To connect to a database from IBM Data Studio, open the **Database Administration** perspective. On the top right corner, click the **Open Perspective** icon and select **Database Administration**.

On the **Administration Explorer**, right-click the white space or under the **New** menu, select **New Connection** to a database. From the *New Connection* window, you see that you can use the IBM Data Studio to connect to different IBM data sources, as well as non-IBM data sources. Select the database manager and enter the necessary connection parameters. [Figure 4.28](#) shows an example.

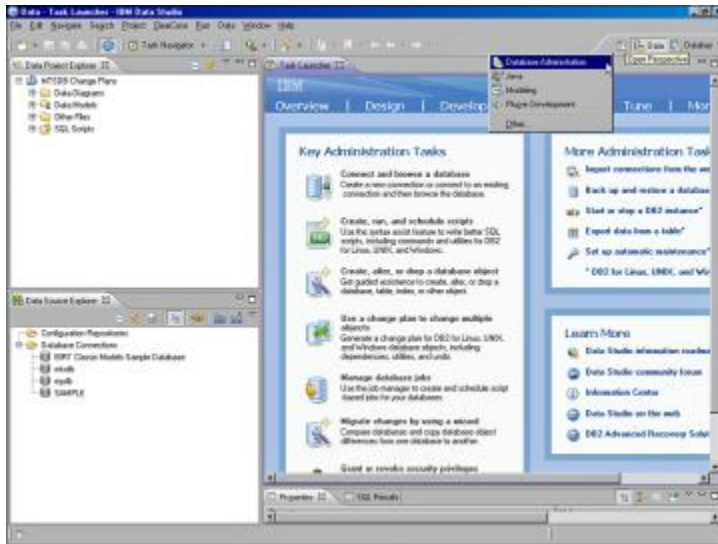


Figure 4.27 Open the Database Administration perspective

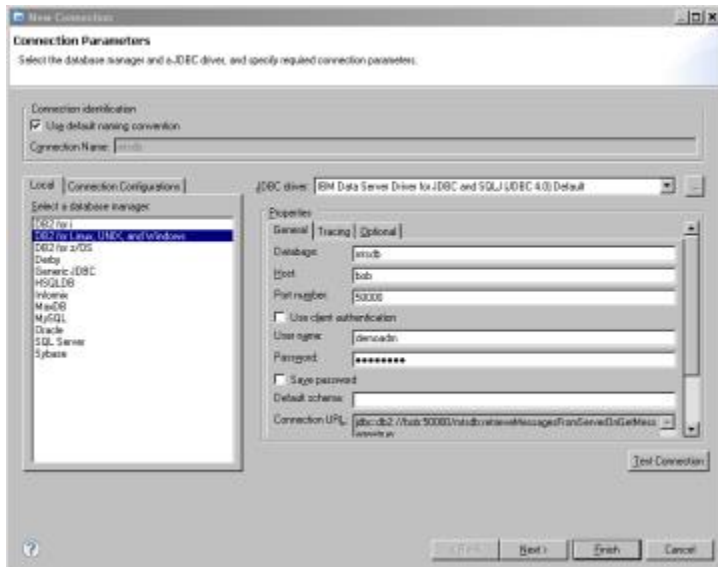


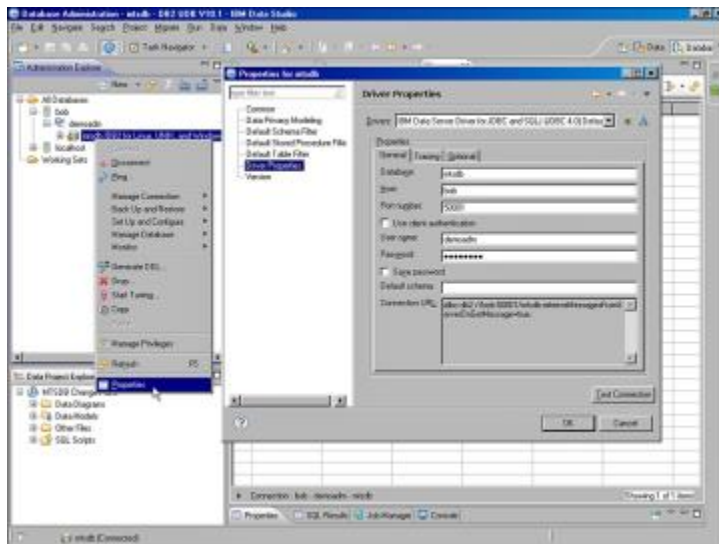
Figure 4.28 Creating a new database connection

Pull down the JDBC driver drop-down menu, and you can select the type of JDBC driver to use. JDBC type 4 driver is used by default.

Use the **Test Connection** button to ensure the connection information you enter is valid. Click **Finish**.

At this point, you have created a connection profile. Connection profiles contain information about how to connect to a database such as indicating the type of authentication to be used when connecting the database, specifying default schema, and configuring tracing options. Other team members can import the connection profiles to their own IBM Data Studio and be able to deploy a set of consistent connection settings.

To update the connection profile, right-click the database and select **Properties**. Properties for the database are displayed as shown in [Figure 4.29](#).



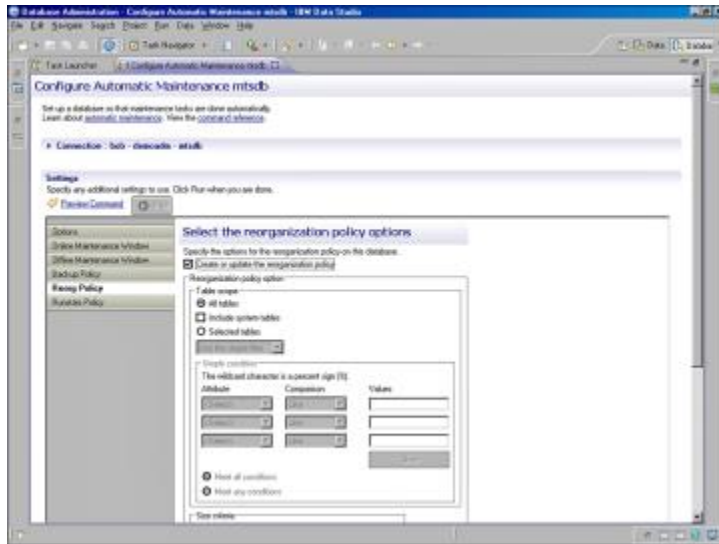
[Figure 4.29](#) Updating the connection profile

General Database Administration Tools

There are few other useful administration tasks available in the menu illustrated in [Figure 4.29](#).

The *Manage Connection* function enables you to rename the connection profile, delete the connection profile, change the user ID and password, and duplicate the profile. The *Back Up and Restore* function enables you to setup a database or table space backups. In the appropriate editor, you can specify the type of backup, location of the backup images, and performance options for the backup. Database backup and recovery is discussed in Chapter 10, "Maintaining, Backing Up, and Recovering Data."

The *Set Up and Configure* function enables you to configure the database. Database configuration and this IBM Data Studio function are covered in detail in Chapter 5. Notice from the menu, you can launch the Configure Automatic Maintenance editor. DB2 provides automatic maintenance capabilities for performing database backups, reorganizing tables and indexes, and updating the database statistics as necessary. The editor enables you customize the automatic maintenance policy (see [Figure 4.30](#)).



[Figure 4.30](#) Select the Automatic Maintenance policy options

The *Manage Database* function enables you to start and stop the database. In DB2, that means activating and deactivating the database. Activating a database allocates all the necessary database memory and services or processes required. Deactivating a database releases the memory and stops DB2 services and processes.

The *Monitor* function launches the IBM Data Studio Web Console. Refer to the section, “IBM Data Studio Web Console,” for introduction of the tool.

The *Generate DDL* function uses the DB2 command-based tool `db2look` to extract the Data Definition Language (DDL) statements for the identified database objects or the entire database. This function and tool come handy when you want to mimic a database, a set of database objects, or the database statistics to another database. As a result of the Generate DDL function in IBM Data Studio or the DB2 command `db2look`, you receive a DDL script. The script contains statements to re-create the database objects you have selected. See [Figure 4.31](#) for a reference of the types of statements you can generate using the IBM Data Studio.



[Figure 4.31](#) Generate DDL function in the IBM Data Studio

For complete options for the DB2 command `db2look`, refer to the DB2 Information Center.

The *Start Tuning* function configures the database to enable query tuning. You might receive a warning indicating that you need to activate the InfoSphere Optim Query Workload Tuner (OQWT) license for advanced tuning capability. Note that IBM DB2 Advanced Enterprise Server Edition comes with OQWT. Follow the instructions to apply the product license or click **Yes** to configure the database server for tuning with the features complementary in the IBM Data Studio.

When the database is configured to use the tuning advisors and tools, you are presented with the Query Tuner Workflow Assistant, as shown in [Figure 4.32](#).

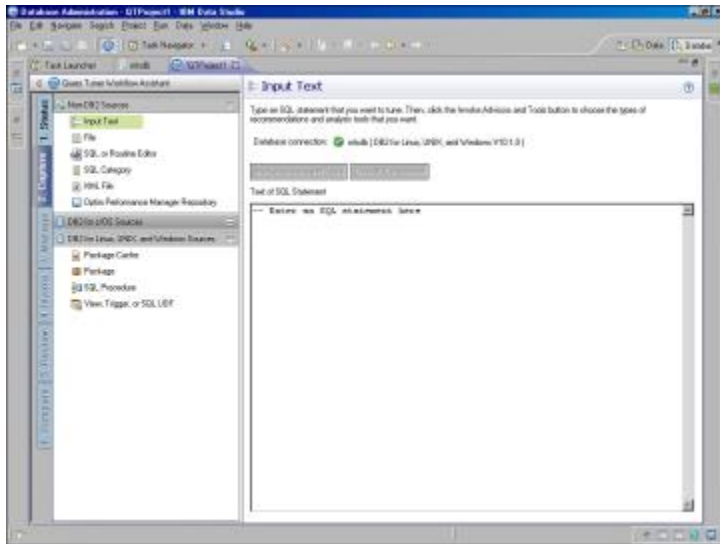


Figure 4.32 The Query Tuner Workflow Assistant

From the *Query Tuner Workflow Assistant*, you can obtain a statement from various sources and tune the statement. In the Capture view, it gives you a list of sources where you can capture the statements. [Figure 4.33](#) shows an example on capturing the SQL statements from the Package Cache. This example captures over 100 statements. Right-click the statement in which you are interested and select **Show SQL statement** or **Run Single-Query Advisors and Tools on the Selected Statement**.

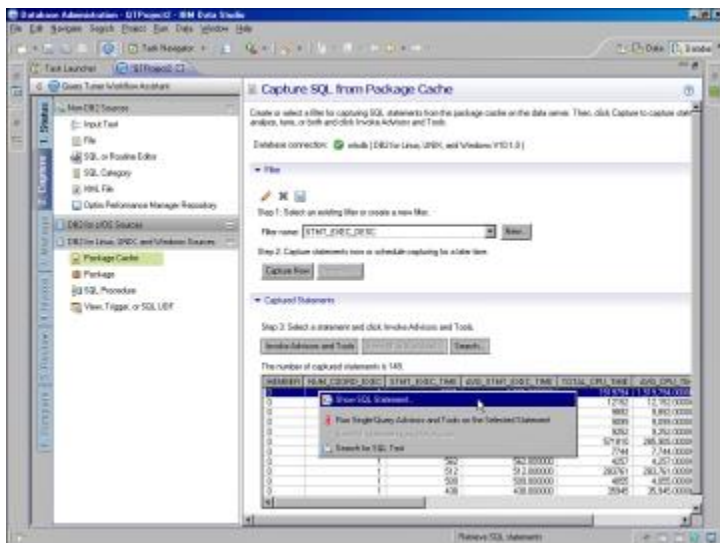


Figure 4.33 Capturing SQL statements for tuning

Run the query advisors and tools on the selected statement. You can now enter the Invoke view. The tool collects information and statistics and generates a data access plan (see [Figure 4.34](#)).



Figure 4.34 Tune query in progress

When the query tuning activities are complete, you are brought to the Review view. It presents you the analysis results and an advisor recommendation, such as the one shown in Figure 4.35. The tool documentation recommends gathering and re-collecting all of relevant statistics of the query.

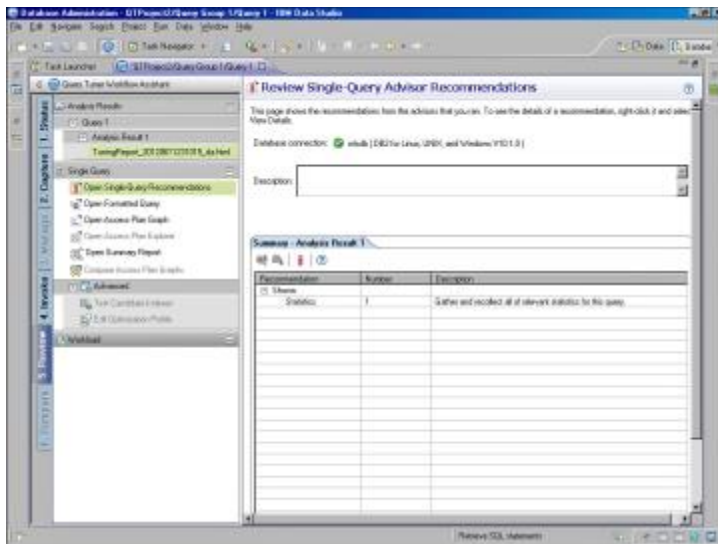


Figure 4.35 Review the advisor recommendation

You can also review the access plan graph generated by the DB2 explain function (see Figure 4.36 for an example). Remember to save the analysis for future references and compare them if needed.

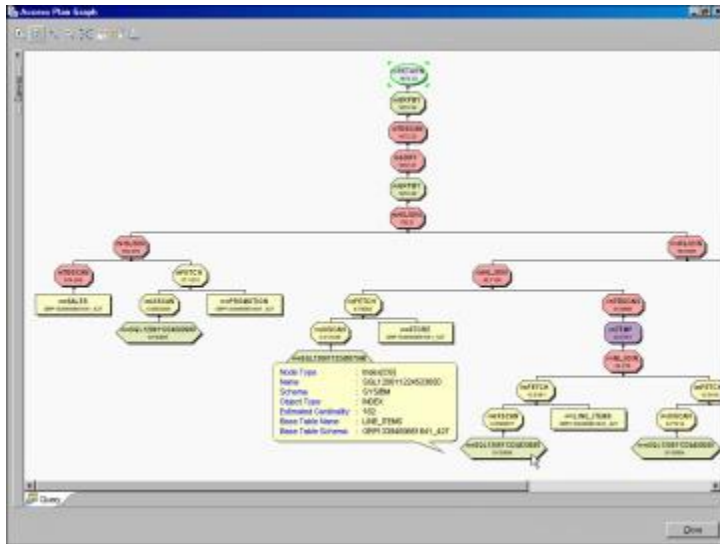


Figure 4.36 Sample access plan graph

The *Manage Privileges* function allows you to grant database privileges to the users. Refer to Chapter 8, “Implementing Security,” for details about privileges and database access controls.

General Database Development Tools

IBM Data Studio consolidates the database administration and database development capabilities. From the *Task Launcher– Develop*, you find a list of key development tasks such as creating and running SQL statements, debugging stored procedures, and user-defined functions (UDFs). Each task brings you to a tool that helps you accomplish it.

SQL and XQuery Editor

The SQL and XQuery editor helps you create and run SQL scripts that contain more than one SQL and XQuery statements. To launch the editor, open the **Data Project Explorer**; under SQL Scripts select **New > SQL or XQuery Script**. As shown in [Figure 4.37](#), a sample SQL script is entered. You can configure the run options for the script.

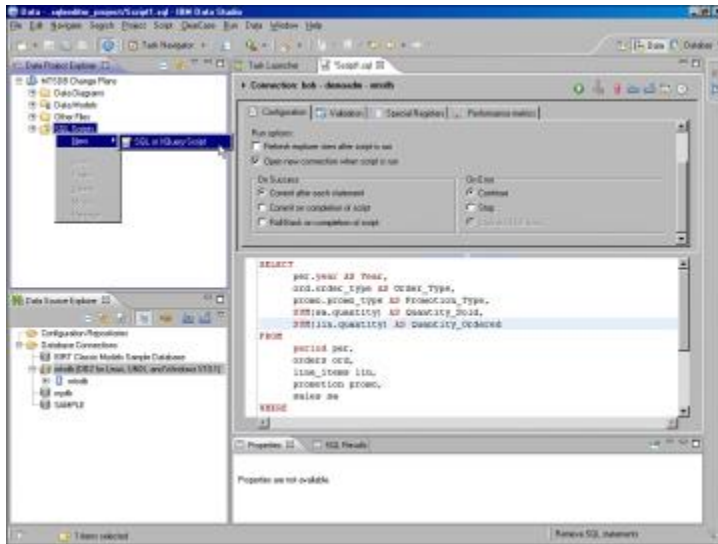


Figure 4.37 SQL and XQuery editor

The editor formats the SQL statements nicely and provides syntax highlights for easier reading as you enter the SQL statements. The functionality content assist is also very useful. It lists all the existing schemas in the database so that you can just select one from the drop-down menu. The editor also parses the statement and validates the statement syntax. You can validate the syntax in scripts with multiple database parsers and run scripts against multiple database connections.

SQL Query Builder

The SQL Query Builder enables you to create a single SQL statement, but it does not support XQuery. As the name implies, the tool helps you build an SQL statement. It helps you look at the underlying database schema or build an expression, as shown in [Figure 4.38](#).

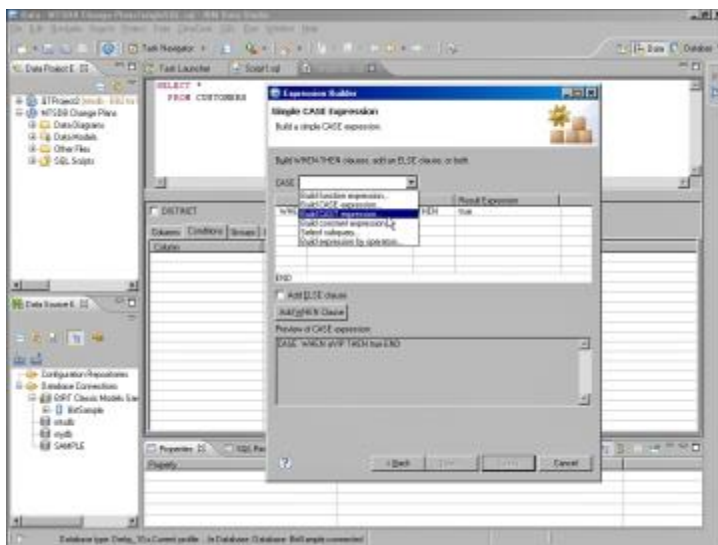


Figure 4.38 SQL Query Builder

NOTE

To get to the SQL Query Builder, click the Data perspective button available at the top right of the window. In the Data Project Explorer view, expand the data profile, right-click the **SQL Scripts** folder, and then click **New > SQL or XQuery Script**. This launches the SQL Query Builder.

Database Routines Editor and Debugger

Stored procedures and user-defined functions (UDFs) are database application objects that encapsulate application logic at the database server rather than in application-level code. Use of application objects help reduce overhead of SQL statements and the results that are passed through the network. Stored procedures and UDFs are also called routines. IBM Data Studio supports routines development and debugging.

From the Data Project Explorer, create a new Data Development Project. In the project, you can create various types of database application objects such as stored procedures and UDFs (see [Figure 4.39](#)). To debug a routine, right-click the routine and select **Debug**.

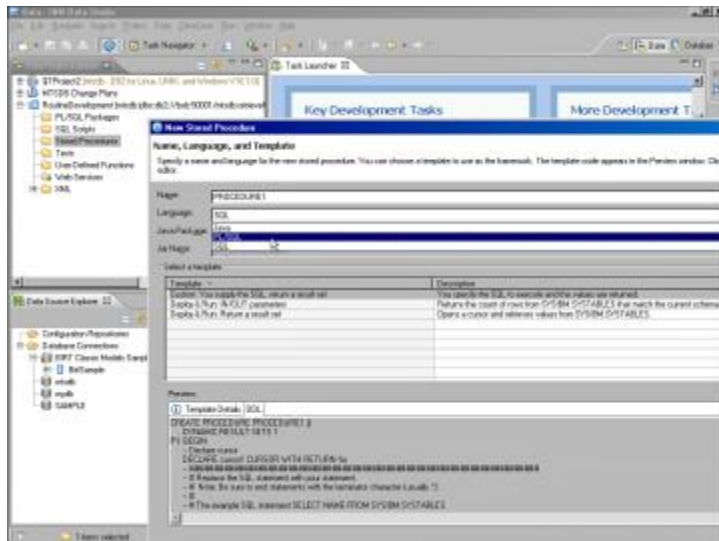


Figure 4.39 Creating a stored procedure

Courtesy: <https://www.informit.com/articles/article.aspx?p=2153657&seqNum=3>

Modified: 2021.10.04.7.15.AM

Dököll Solutions, Inc.

