

Tech Tip: How'd You Get Your SQL Server Data in My IBM i DB2 Data!?! ---

IBM DB2 Web Query integrates with Microsoft SQL Server.

Published Friday, 05 November 2010 00:00 by MC Press On-line [Reprinted with permission from iTechnology Manager, published by MC Press, LP; <http://www.mcpressonline.com>.]

Written by Robert Andrews – robert.andrews@us.ibm.com

DB2 Web Query is a great tool for reporting and graphing data on IBM i. It has built-in capabilities that allow DB2 Web Query to access data on any IBM i Power System. This starts to bring the power of a federated database into IBM i. A federated database system is one in which a single system accesses data that resides on multiple database systems. In the past, data was often replicated from one system to another to allow this type of access. However, this leads to issues with data being out of date, out of sync, and out of control. To help extend federated database capabilities, DB2 Web Query introduced the Microsoft SQL Server adapter.

Microsoft SQL Server Adapter

Available as an additional feature to the base DB2 Web Query product, the Microsoft SQL Server adapter allows DB2 Web Query to reach out and pull data from any SQL Server version 2000, 2005, or 2008. Once the metadata is defined to link the SQL Server data into DB2 Web Query, the tools work exactly the same. This allows a report developer to quickly use the same familiar toolset to build powerful reports over IBM i data or Microsoft SQL Server data!

This capability also allows for the integration of data from both platforms within a single report! Imagine joining your manufacturing data on IBM i with your Point of Sale data on SQL Server within one report. And all the data is accessed live, in real time, with no replication or additional administration tasks.

Setup

There are two steps to set up the Microsoft SQL Server Adapter for DB2 Web Query. The first step is to install and configure the adapter. Once the adapter is set up, metadata to the SQL Server box needs to be created.

To install the Microsoft SQL Server, insert the 5733QU2 DVD and install Option 6. Next, download the correct JDBC driver from Microsoft's Web site, either version 1.1 for SQL Server 2000 or 1.2 for SQL Server 2005 and 2008. Place the JAR files into the /QIBM/UserData/Java400/ext directory.

Log in to the DB2 Web Query interface and launch the metadata tool. Select Add Adapter, and choose the correct version. Set the correct URL and driver strings. For SQL Server 2000, use `jdbc:microsoft:sqlserver://sqlserver.mycompany.com:1234` and `com.microsoft.jdbc.sqlserver.SQLServerDriver`. For SQL Server 2005 and 2008, use `jdbc:sqlserver://sqlserver.mycompany.com:1234` and `com.microsoft.sqlserver.jdbc.SQLServerDriver`. Of course, replace `sqlserver.mycompany.com` with the SQL Server's DNS name or IP address and 1234 with the correct port for the SQL Server.

For security, the SQL Server must be set up for both Windows and SQL Server authentication. The adapter will pass a single, static user profile and password to the SQL Server. This profile will need to be authorized to read all the data required to be accessed by the adapter. End and restart DB2 Web Query to pick up these changes.

Now that the adapter is set up, create metadata against the resources that need to be accessed by the SQL Server. Just as you did with the DB2 CLI adapter, right-click the SQL Server adapter and select Create Synonym. The only difference here is the addition of a dropdown box to select the database on the SQL Server. Unlike DB2 on IBM i, SQL Server breaks data into multiple databases within a single SQL Server and then into schemas and tables. In sticking to IBM's best practices for DB2 Web Query, IBM would recommend using a suffix of something like `_sqlsrvr` as a visual queue that this data is remote.

And that's it! The synonyms to the SQL Server data show up right along with the local data synonyms. They function and operate exactly the same way. This is because of the power of the metadata layer. As report builders, the location of the data is not important! Build reports using the same, common toolset, regardless of the source data's location.

For complete, illustrated details on the installation and configuration, please see chapters 14.1 to 14.3 of the [Getting Started with DB2 Web Query for i Redbook](#).

Performance

While the capabilities exist to perform these cross-system joins, keep in mind that performance may vary from that of a local query. When a join is done, the system attempts to eliminate as many rows as possible as soon as possible. The system uses a variety of tools, such as indexes and statistics, to accomplish this. When DB2 Web Query performs a cross-system join, the same techniques cannot be used.

This may result in larger result sets coming back into DB2 Web Query, which then processes the join, eliminating data at that point. Keep in mind that your final result will be accurate, but the overall time required to execute these cross-system joins is longer than if the data were all local.

In addition to data processing times, the speed of the network and the relative distance between the servers will also affect the overall run time. If the Power Systems and Microsoft servers are in the same data center, the transmission times will be quicker than if one is in an East Coast data center and the other in a West Coast data center.

More Information

If you would like more details about the Microsoft SQL Server Adapter for DB2 Web Query, please visit the [DB2 Web Query site](#). For pricing or ordering, please contact your Business Partner or IBM marketing rep. The best way to learn DB2 Web Query is to follow the tutorial-driven IBM Redbook [Getting Started with DB2 Web Query for i](#).