JACKIE's Forum

Is DB2 Really DB2?

Jackie Jansen

The short answer is YES. Now let me clarify. I often get asked about DB2 UDB running on an i5 / iSeries / AS/400 systems and how that database compares with DB2 running on other platforms.

IBM has three flavours of DB2. We have DB2 UDB for zSeries, DB2 UDB for i5 and DB2 UDB for LUW. The latter stands for Linux/Unix/Windows and is sometimes referred to as DB2 for Distributed. Each member of the DB2 family has a different code base and is developed in a different laboratory. This is largely to enable the database to take full advantage of the appropriate operating system and hardware platform where that makes sense. All three development labs and our research lab are represented on IBM's internal DB2 council.

The majority of questions that I get asked are from the viewpoint of source code compatibility or skills transfer capabilities. Questions like: Does my training on DB2 UDB for zSeries translate to DB2 UDB for i5? Can I easily port my application from DB2 UDB for LUW to my System i5?

From a user's point of view there are three basic areas to consider. The first two areas are comprised of SQL programming interfaces. SQL consists of Data Definition Language (DDL) and Data Manipulation Language (DML) specifications. DDL is used to create and manage objects such as tables, views and indexes. The Data Manipulation Language or DML is used to access the data. SELECT, UPDATE, INSERT and DELETE are all examples of DML statements.

DML across the three family members has a very high degree of compatibility. DML source code is over 95% compatible with the more common functionality being closer to 100% compatibility. There are some differences between the systems. Sometimes this is simply the timing of various announcements. i5 chose to bring out support for binary large objects later than DB2 for distributed. This was largely due to the fact that very few of our customers were asking for that capability. On the other hand, from a functionality perspective, DB2 UDB for i5's cursor support is more robust than the support in DB2 UDB for LUW. V5R4 added more functionality to the database on System i5 to increase compatibility with the other DB2 family members. Examples of this new functionality included more support for recursion and additional support for OLAP type capabilities including RANK and ROW_NUMBER.

AS of V5R4 DB2 UDB for i5 is currently the only database, inside or outside the DB2 family, that meets all the SQL 2003 Core Standards.

DDL or Data Definition Language is not quite as compatible as DML. This is largely because defining the database requires a closer tie to the underlying operating system and hardware. For example, when creating tables in DB2 UDB for LUW or zSeries you also define table spaces, buffer pools etc. The i5 or iSeries automatically handles all this storage allocation for you. Between i5/OS and the hardware the storage is allocated for a table and then the data is striped across multiple physical disk arms for maximum performance.

There are some differences in data types between the different DB2s although these are becoming fewer and fewer. One example would be the fact that DB2 UDB for i5 internally stores a NUMERIC field as a zoned decimal. A NUMERIC field would be stored as a packed decimal on UNIX. Variable length columns are handled differently internally. This is more of a performance issue than a straight compatibility issue. For best performance on a System i5 any field less than 40 characters shuld be defined as a CHAR data type. For larger fields you should use the ALLOCATE keyword and reserve a length that will fit most of the values in the column in the base table. Anything larger than the length you have allocated will be stored in an overflow area.

The main differences are not seen in the actual SQL code that you would port between systems but in the administration of the database. On the System i5 we use i5 or iSeries Navigator as our graphical interface into the database. This is tightly coupled with all administration and operational aspects of the System i5 including database performance monitors, visual explain, and operating system functionality such as work management. DB2 UDB for LUW uses DB2 Control Centre for administration. While Control Center has support for the i5 it is not as full function as i5 Navigator support.

In summary, while the databases are not identical they have a lot more similarities than differences. Over the years the documentation on the SQL differences between the DB2 family has gone from a large manual to a small paper. As always, there is additional information at www.ibm.com/iseries/db2. TUG

***Jackie Jansen** is a Senior Consulting IT Specialist. She currently works in the IBM Americas Advanced Technical Support Solutions Centre. Jackie is a frequent speaker at iSeries Technical Conferences andUser Group meetings. Contact her at jjansen@ca.ibm.com.*