



Craig S. Mullins

[Return to Home Page](#)

February / March 2007



zData Perspectives

by Craig S. Mullins

[DB2 9 for z/OS: No V for z](#)

The beta program for the next version of DB2 for z/OS began back in June of 2006. This latest and greatest version of DB2 is known as DB2 9 – no V, just 9. At the time I’m writing this (early December), no GA date has been announced, but “word on the street” says to expect GA early in 2007.

So, what can you expect in terms of new features from DB2 9? Well, the big one is called pureXML. Basically, pureXML allows you to store data as native XML. What’s the big deal? We can already store

XML in DB2 prior to V9 (I know, but I can't just bring myself to say "prior to 9" without the "V")?

DB2 9 changes the XML game. You will be able to search and analyze structured data in a relational data repository and unstructured data in an XML repository without the need to reformat it. The approach is novel in that it will support native XML, basically enabling dual storage engines. When you want to store XML in DB2 9 you no longer have to store it as a CLOB or shred it into tables. Think of XML as just another data type. If you are a big XML user, you'll want to seriously consider storing it in DB2 9.

What else is new? DB2 9 expands support for online schema changes... and IBM renames it to Database Definition On Demand (DDOD). It seems like IBM is constantly renaming things. One new DDOD capability enables you to replace one table quickly with another using cloning. Others allow you to rename columns and indexes.

Online table space reorganization is also significantly improved. As you probably know, when reorganizing just a couple of partitions in a partitioned table space the BUILD2 phase can take a long time to complete. V8 removed the outage for DPSIs, and now DB2 9 removes the BUILD2 phase for all types of secondary indexes.

DB2 9 introduces a new type of table space that combines the attributes of segmented and partitioned. When utilized, DB2 automatically adds partitions as needed to support your rapidly-

growing data. And DB2 9 begins phasing out simple table spaces – you will no longer be able to create simple table spaces.

There are also a bevy of new SQL capabilities. First up, DB2 9 offers new data types for DECFLOAT (decimal floating point), BIGINT (8 byte integer), BINARY and VARBINARY types.

On the query front, DB2 9 delivers support for INTERSECT and EXCEPT in SQL SELECT statements. These two set operations can be used to simplify some SQL statements. Both are similar to UNION: INTERSECT is used to match result sets whereas EXCEPT combines non-matching rows from two results. You might be familiar with other DBMS products that refer to EXCEPT as the MINUS operation.

In other SQL news, DB2 9 allows:

- ORDER BY and FETCH FIRST clauses to be specified at the fullselect level instead of only at the SELECT statement level;
- A new MERGE statement that takes two tables and merges the data (rows in the target that match the source are updated and rows that do not exist in the target are inserted);
- A new TRUNCATE statement (a quick way to DELETE all data from a table);
- the ability to SELECT from DELETE, UPDATE, and MERGE statements;
- INSTEAD OF triggers (the trigger code is run instead of the code that triggered it);
- more OLAP extensions (RANK, DENSE_RANK and ROW_NUMBER);
- and enriched text handling with functions like caseless comparisons, cultural sort, and index on expressions.

Additionally, DB2 9 delivers improved security and regulatory compliance features. Role-based security authorization offers a more flexible technique for controlling authorization than groups or users. A role is a set of users sharing the same security privileges. A user belonging to a particular role can perform the set of tasks and accesses for which permissions are granted for that role. Other security improvements include trusted security context, SSL with encryption on the wire, and improved audit filtering, including the ability to include or exclude based on userid, workstation, application name, package location collection or name, connection id, correlation id, and role.

Finally, an overview of a new DB2 release is incomplete without discussing performance. DB2 9 delivers many new performance-related features. INSERT performance should improve dramatically due to a wide range of improvements including the ability index on expressions and randomized index keys. Logging performance is improved substantially and larger index page size options are provided to reduce the number of page splits.

Several improvements have been made to DB2 optimization, including a histogram for DB2 statistics so that DB2 can gather a better understanding of skewed data over ranges. The DB2 optimizer can deploy cross query block optimization instead of having to work on one query block at a time.

Visual EXPLAIN has been improved into an Optimization Support Center with improved instrumentation for analyzing and improving query performance.

Of course, this is just a quick overview of DB2 9 – there are many more features than I can adequately cover in a column of this nature. Use it as a starting point for your DB2 9 learning curve and be sure to keep your eyes open for additional details coming from IBM.

From [zJournal](#), Feb / Mar 2007

© 2007 Craig S. Mullins, All rights reserved.

[Home](#).