IBM DATABASE MAGAZINE » Powering business information

Quarter 1, 2007 Vol. 12, Issue 1

http://www.dbmag.intelligententerprise.com/iduguserview/showArticle.jhtml;jsessionid=NFZBSDVE5C3OHQE1GHOSKHWATMY32JVN?articleID=199201845

Boosting DB2 9 for z/OS Performance

By David Beulke

New features produce dramatic application results.

DB2 9 for z/OS has many new features that give database and application designers greater flexibility to solve business problems. I recently helped a client with an application design review and saw first hand how implementing many of these features can immediately and dramatically improve performance.

The first improvement the client implemented was the `SELECT FROM UPDATE/DELETE`. This capability extends the functionality introduced with `SELECT FROM INSERT` in DB2 Version 8 to these additional SQL statements. My client generates new key values that need to be returned to the application during an UPDATE statement; now, this task can be handled with a one SQL call to the database instead of two.

DB2 9 also introduces the new SQL statement `MERGE`. `MERGE` improves application performance by allowing DB2 to integrate data into a table in a single SQL statement. The `MERGE` statement could improve my client's nightly batch database `INSERT` and `UPDATE` process by allowing DB2 to avoid doing the `SELECT` data existence check and then an `INSERT` or `UPDATE` to the table. Now a SQL `MERGE` statement will determine whether the database-keyed row is already in the table and will integrate the data into the database.

The `MERGE` SQL statement operation can use multiple rows as an input array. When multiple input rows are used, the phrase `NOT ATOMIC CONTINUE ON SQL EXCEPTION` can be specified. This phrase allows DB2 to process each row independently. If an error occurs when merging a row, only the error row is backed out. DB2 continues processing all the subsequent input rows. DB2 will take each of the multiple input rows then determine the correct insert or update operations to integrate them into the table. When there are triggers over the table, each successfully merged row fires the trigger and the appropriate actions are taken. MERGE error rows encountered don't fire triggers.

Applications can use the new `MERGE` SQL statement in a `SELECT FROM MERGE` statement. With only one SQL statement, DB2 can insert or update multiple rows in the table and return the values back to the application. This capability reduces application complexity and helps eliminate potentially millions of SQL calls in nightly OLTP or data warehousing system load and update processing cycles.

Another great new application feature is the ability to use an `SQL ORDER BY` or `FETCH FIRST n ROWS` within a subselect and fullselect. This feature gives applications the ability to quickly retrieve a small number of rows in the proper order. Using `ORDER BY` affects which rows are first in the answer set; the `FETCH FIRST n ROWS` clause helps limit the number of rows retrieved.

My client's OLTP and data warehousing applications only need to retrieve top products or one screen of data, so this feature fits their needs. Retrieving a limited number of rows improves performance and reduces coding complexity for all applications.

Other great DB2 9 for z/OS application features include new scalar functions, global query optimization, autonomic re-optimization, and many more. Take advantage of all the early customer user experiences presented at the IDUG North America conference this May in San Jose, Calif. This conference offers perspectives presented by users, consultants, IBMers, and third-party vendors.

***David Beulke**, a past president of IDUG, is a DB2 consultant, author, and lecturer who specializes in database performance, data warehouses, and Internet applications.*

Return to Article