

DB2 Version 3: A Closer Look

By Craig S. Mullins

DB2 Version 3 was announced by IBM this past March and there have been countless articles in trade journals and presentations given by industry experts on the new features of this latest and greatest DB2. Yet simply knowing "just the facts" will not provide you with the information necessary to successfully plan for, and implement, DB2 V3. What is needed—in addition to the facts—is an analytical and critical view of these features. What will I need to change? How should this impact my budget? Do I need to re-evaluate potentially obsolete ways of handling my DBA tasks? In short, how will DB2 V3 impact my world?

This article will provide a quick look of some of the new features provided by DB2 V3. The approach taken will not be descriptive, but analytical. Hopefully this will provide you with a different insight into the new DB2.

Data Compression

With each successive release of DB2, IBM is providing more and more options for compressing DB2 data. As DB2 is stretched to its maximum capacity the need for quick and reliable compression becomes more and more obvious.

DB2 V2.3 provided an EDITPROC for compression (DSN8HUFF), but it was not competitive with third party compression tools such as PLATINUM Data Compressor. DB2 V3 goes even further by embedding a data compression feature into the "guts" of DB2. Compression is indicated in the DDL by specifying COMPRESS YES. Likewise, it must be turned off in the DDL by specifying COMPRESS NO. DB2 V3 compression provides two very clear benefits: (1) hardware-assist compression and (2) the ability to store 255 rows/page instead of the previous limitation of 127 rows/page.

Let's examine the hardware-assisted compression a bit closer. It is similar in function to the hardware-assisted sort that was offered for DB2 V2.3. Hardware compression is available only to those users owning IBM's high-end CPU models (ES/9000 Model 511 or 711). This does not mean that the new DB2 V3 compression features are only available to those with high-end CPUs. Hardware-assisted compression simply speeds up the compression and decompression of data—it is not a requirement for the new data compression features of V3.

What is the impact of the new compression features? Overall, the addition of the new data compression technique within DB2 can be seen as a good thing for most users. Some who never looked at compression before will re-evaluate their compression needs. But how will it impact third party tools? Before refusing to evaluate third party tools consider the following:

- IBM compression supplies only a single compression routine (based on the Ziv-Lempel algorithm) whereas PLATINUM Data Compressor provides several different compression routines. This enables the user to better fit the algorithm to the composition of the data—using different compression algorithms for different types of data.
- It is possible to achieve 255 rows/page even without using IBM's compression. Simply alter the tablespace to indicate COMPRESS YES and then alter it right back

specifying COMPRESS NO. DB2 modifies the DBD to indicate that 255 row pages are permitted when the tablespace is altered to indicate COMPRESS YES. However, DB2 does not modify the DBD after the second alter to set the tablespace back to COMPRESS NO status. If the database is ever dropped and re-created to perform maintenance you will need to re-issue the alter statements to ensure 255 rows/page.

- Third party tool vendors are constantly enhancing their products to take better advantage of the operating system and the hardware environment. To ensure that you are getting the best bang for your buck in terms of data compression, it is wise to evaluate all of your options before settling on any given one.

Bufferpools

I have listened to many an old IMS DBA grumble about the limited number of bufferpools available to DB2. Up to and including DB2 V2.3 DBAs were constrained to having only four bufferpools for DB2: three 4K pools and one 32K pool. Well, with DB2 V3 there are now 60 of 'em. What in the world are we gonna do with 60 DB2 bufferpools? Plenty!

For starters, it should be much easier to isolate performance-critical tablespaces and indexes into their own bufferpool. Why would you want to do this? Consider the situation whereby several large tablespaces are assigned to the same bufferpool. A sequential prefetch request for one of these could monopolize up to 80% of the bufferpool! This would mean that most of the other tablespaces and indexes assigned to this bufferpool would be flushed out causing costly disk access to be incurred the next time any of that data needs to be accessed.

Here are a several more ideas for the utilization of the new bufferpools:

- Isolate objects by functional group. For example, you could isolate the DB2 Catalog in its own bufferpool (note: this must be BP0). Likewise, you could isolate the QMF administrative tables, QMF tablespaces used for SAVE DATA, and objects used by third party tools.
- Create application-specific bufferpools. Each application could have its own bufferpool, reserved for objects used by that application only.
- Dedicate one of the bufferpools to DSNDB07 to improve the efficiency of DB2 sorting.
- Place active code tables into their own bufferpool, just large enough to hold the entire table.
- Separate tablespaces from their associated indexes to achieve a higher buffer hit ratio for index pages.
- Separate objects based upon type of access. For example, separate ad hoc objects from transaction-oriented objects; or sequentially accessed objects from those which are randomly accessed.
- Reserve one of the bufferpools for tuning and problem solving. By altering specific objects to the tuning bufferpool, I/O can be isolated for those objects in a single pool enabling easier performance monitoring.

Be careful as you implement additional bufferpools. The more you implement, the more system memory will be required. Be sure to analyze the affects that additional bufferpools will have upon your system. You want to enhance I/O processing, but avoid paging. And don't be stingy with the memory for DB2. A bufferpool set at 1000 provides

4 megabytes of storage. Most desktop PCs come factory-equipped with that much memory! If your PC needs that much memory, imagine what DB2 could use!

One final caution regarding the bufferpool enhancements for DB2 V3: remember that BP32 and BP32K are two different animals. BP32 is one of the fifty 4K bufferpools available with DB2 V3. BP32K is one of the ten 32K bufferpools. If you miss, or add, an erroneous "K" you may wind up using or allocating the wrong bufferpool.

Synopsis

This article has analyzed only two of the numerous features that will be available with DB2 Version 3. Future SYSJOURNAL articles will take a look at some of the other new features, but until then remember this: DB2 V3 will provide better overall performance but at the expense of more memory, upgraded hardware, a new approach to the evaluation of data compression options, and more complicated bufferpool management.