IBM Information Management software

Storage optimization in DB2 9

Highlights

- ***Eases the strain on storage resources by reducing the size of table data***
- ***Includes a compression estimator feature to help you determine potential savings***
- ***Minimizes I/O and uses memory more effectively***
- ***Maximizes compression capabilities by looking at data patterns in an entire table resulting in up to 80% storage savings***

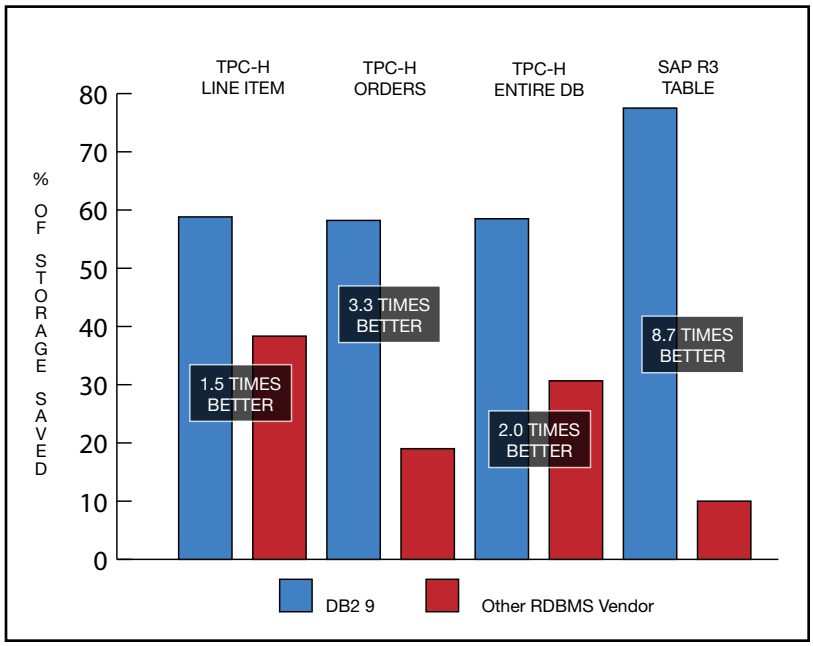

The breakthrough compression capabilities within DB2 9 can help you realize significant cost savings.

Data — both structured and unstructured — continues to grow at a staggering rate. Along with the ever-increasing volumes of data comes the demand for more storage. Preserving the integrity of your data is first and foremost, but keeping costs in check is also critical.

The latest version of DB2® for Linux®, UNIX® and Microsoft® Windows®, DB2 9 formerly code named Viper, introduces leading-edge technology that enables you to address the continued strains on your available storage resources through powerful new data compression features. The end result: Less of your budget spent on storage and the potential to significantly increase performance.

The next step forward

Storage optimization within DB2 9 is based on a compression methodology that maps repeating data patterns within a table to symbols of a smaller size with the ultimate goal of reducing the size of the table data. DB2 first enabled compression capabilities in Version 8.1, which allowed users to efficiently compress and store null values within a system. Version 8.2 added backup compression capabilities, which lets users compress data before it is backed up.

A significant difference

"With the new compression technology in DB2 Viper, we realized an 80% improvement in space savings for our most critical tables in our Data Warehouse," said Donny Ledbetter, Senior Database Administrator, AutoZone. "We were even more pleased with this technology when we found that Viper's compression capability helped us process queries to the database an average of 40% faster than before. We're looking forward to seeing the same results with our Operational Data Store and OLTP systems."

The release of DB2 9 brings even more features into the compression portfolio, including:

Row-level compression. Modeled after DB2 for z/OS®, row-level compression looks for repeating patterns in your table and builds a dictionary based on the table values. Row-level compression can compress repeating patterns within a column or patterns that span multiple columns resulting in very high compression ratios. It is completely transparent to your application. This approach differs from other offerings that build a compression dictionary for each page. By building a compression dictionary at the table rather than the page level, DB2 9 analyzes the entire table to identify patterns resulting in significantly better compression and more savings on storage.

DB2 INSPECT facility. This feature allows you to estimate the potential compression savings by calculating the number of pages saved on a given table.

Decrease your total cost of ownership

Savings on disk space, although appreciable, is only one aspect of the savings that compression offers. Many companies replicate their production database at least two times over for test and development environments. It is also not uncommon for companies to store at least one backup — if not more — on disk. When you consider that any disk savings you see with your data is multiplied through backups, these replicated versions of the database can quickly add up to exponential levels of savings. And when you take into account the savings on rack space, cables, floor space, UPS and all the IT infrastructure that supports storage equipment, the potential savings are even more dramatic.

Increase the performance of your overall system

Beyond storage infrastructure savings — which can be substantial — compression can also translate to higher performance. Many systems are I/O bound, which is necessarily going to be the slowest mechanical system within your overall database system.

Using row compression reduces the number of I/Os required to retrieve data and can thereby help reduce I/O bottlenecks to improve overall performance. Even on systems that are CPU bound, compression can still improve performance. The tradeoffs of CPU cycles vs. I/O bandwidth for many workloads is a positive tradeoff, with many customers reporting significant improvements in system performance. To uncompress data at runtime, DB2 9 will use more CPU cycles. However, since rows remain compressed in the buffer-pool, DB2 9 can still improve performance by making more efficient use of available memory and by processing data faster.

For more information

Part of a wide array of information management solutions, DB2 9 provides unmatched development and runtime efficiencies, performance advantages and reduced costs. To learn more about how DB2 9 compression features can help you control the escalating costs of storage management, contact your local IBM Business Partner or visit: **ibm.com**/db2/viper/

IBM Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
06-06

G507-0715-00