

Reduce Data Storage Costs and Increase Performance in SAP Environments

A Deep Look at IBM DB2 9 Deep Compression



Cong Lin
Senior Software Engineer, IBM

Cong Lin is a member of the IBM SAP Integration and Support Center at the IBM Toronto Lab. Her current activities include testing SAP applications with DB2 for LUW and helping customers with problem analysis and troubleshooting. She is also a customer advocate, providing support for large customer accounts running SAP and DB2 for LUW.



Marianne Nesiem
Senior Software Developer, IBM

Marianne Nesiem is a member of the IBM SAP Integration and Support Center at the IBM Toronto Lab. Her current activities include testing and certification of SAP R/3 with every new release of DB2 UDB, and helping customers running SAP in a DB2 UDB environment with problem analysis and troubleshooting. She is also a customer advocate, providing custom-tailored support for large customer accounts running SAP and DB2.

Organizations are generating increasing amounts of data, and they're keeping their existing data longer — as a result, storage costs are on the rise. With more disk space being consumed, companies need to get innovative to reduce the TCO of storage.

For those companies dealing with massive amounts of data in SAP environments, IBM DB2 9 Deep Compression not only reduces storage costs, but also enhances performance. IBM customers have experienced data storage reductions of up to 70% and SAP response time improvements of more than 30% by implementing DB2 9 Deep Compression.

The Benefits of DB2 9 Deep Compression in SAP Environments

Jointly developed by IBM and SAP, DB2 9 Optimized for SAP solutions has numerous built-in features customized to improve performance and resource utilization. Deep Compression is one such feature.

The results reported by customers using this solution have been great. A large, global SAP customer using SAP NetWeaver Business Warehouse (SAP NetWeaver BW) and IBM DB2 9.5 Deep Compression lowered its 16TB database storage by 50%. Deep Compression also lowered the number of database pages by 48% and improved query response times by 23%. Here, compression had no impact on batch performance and the customer's most important online transactions became 20% faster with IBM DB2 9.

DB2 9 Deep Compression: Features and Benefits

Let's take a closer look at three specific features of DB2 9 Deep Compression, namely automatic storage, "large object" (LOB) in-lining, and index compression, all aimed at improving performance and lowering administration costs.



Automatic Storage

Automatic storage makes storage management easier. Rather than managing storage at the table-space level using explicit container definitions, database managers can manage storage at the database level, meaning that they can take responsibility for creating, extending, and adding containers.

To take full advantage of this feature, database administrators must enable automatic storage at both the database level and the table-space level. You can enable automatic storage at the table-space level in an online or offline format. The online method is cheaper than taking the database offline, but using the online format involves moving data from non-automatic storage containers to the new automatic storage containers. Since this rebalancing will move all data in the table space, it will affect system performance. That's why it is important to perform this rebalancing only when the system is not running peak loads. The offline method doesn't require you to rebalance the table spaces since the containers that were created at the time of restore automatically manage the data.

Large Object In-Lining

An LOB is a data type that stores large data. LOB values are usually much larger than other values in a table and are generally stored separately from the table that references them. In general, LOB data is not eligible for row compression. However, with DB2 9.7, it is now possible to include smaller LOBs in the base table row, compressing the data and making it easier to manipulate data and eliminate access overhead. This operation, referred to as in-lining, improves the performance of queries that access LOB data, because no additional I/O is required to fetch, insert, or update the data.

In-lining has shown great results in SAP environments, especially in conjunction with DB2 9's Deep Compression feature. In a case study involving over 250 SAP tables that contained LOBs, we found that compressing tables with Deep Compression before or after in-lining could make a significant difference in the compression ratios. For instance, one SAP table with LOBs could be compressed by 20% using Deep Compression. However, when we performed in-lining and then compressed the table, we were able to compress the table by 93%. Similarly, the space savings with

Deep Compression on another SAP table with LOBs went up from 7% to 82% after in-lining.

In addition, we found that while some tables with smaller LOBs did not notice a significant decrease in total size, there was still an improvement in performance since the data was moved to the base table and it became more quickly accessible.

Index Compression

In many SAP environments, it is common to have very large indexes that are defined on large tables. As is the case with data objects, Deep Compression compresses index objects to further reduce disk consumption and storage cost with minimal effect on system performance (in most cases).

To minimize storage space, database managers can use multiple compression algorithms, which are fixed so that the process requires no dictionary to compress indexes. As of DB2 9.7, database managers can also compress temporary tables — as well as indexes on those tables.

This compression can be done online or offline. The degree of compression will vary based on the type of index, as well as the data the index contains. For example, in an index where there is a high degree of commonality in the prefixes of the index keys, a database manager can apply compression based on the similarities in prefixes.

Learn More

IBM DB2 9 offers customers numerous features and functionalities that help reduce overall operating costs. The Deep Compression feature, in particular, helps SAP customers handle large amounts of data, enabling them to compress the data to lower storage costs and improve system performance.

For more information on IBM DB2, visit www.ibm.com/db2. ■

Quick Tip

IBM and SAP have jointly included an option within the SAP DBA Cockpit that directly invokes compression on SAP tables. In addition, an SAP-provided report identifies a list of the best candidates for compression, making DB2 9 Deep Compression not only important for storage savings, but also easy to implement in SAP environments. The report is available at <https://service.sap.com/sap/support/notes/980067>.