Analytics in Action with Teradata® In-Memory Optimizations

Performance Study by Large Manufacturer

Richard Hackathorn, Bolder Technology

03.16 EB9292

Table of Contents

Enterprise data warehouse is the most critical component of their IT architecture.

This study discusses a customer's experiences with in-memory optimization features contained in Teradata® Database 15.10, along with the business and technical impacts of those features. Database technology has matured greatly over the past four decades. The latest database engines are leveraging memory for increased speed and flexibility. This study focuses on the Teradata Database's specialized techniques for optimizing data flow through memory to/from the CPU, improving system performance and thus user experience across the corporation.

Context

This study documents the experience of a large corporation exploring the business implications of low-level in-memory optimizations to its enterprise data warehouse. Such low-level hardware issues do not usually concern business executives. However, in-memory optimizations are having major performance improvements that enable major business benefits. As the following case study illustrates, executive should be aware of the cost savings and potential business impacts.

For most corporations, the database engine for the enterprise data warehouse is the most critical component of their IT architecture. To perform essential business functions, it must operate efficiently and reliably every second of every business day. For data-warehousing (DW) vendors, like Teradata Corporation, the maturity and evolution of the database engine is a key factor in retaining market share, since corporations are always pushing the capabilities of their data warehouses.

Those data warehouses typically contain close to a petabyte of critical business data, organized into hundreds of tables, used by many thousands of persons, performing tasks across all business functions. Further, the business opportunities of business analytics and social media data are pushing a single-platform data warehouse into a complex ecosystem of platforms, spread across on-premises datacenters and hybrid cloud architectures.

At the heart of these complex data ecosystems is the low-level hardware mechanism for processing each small chuck of data through memory to/from the CPU. The complexities of current computer architectures challenge DW vendors to organize data in memory so that it flows efficiently through this execution pipeline. Hence, an obscure hardware mechanism can significantly affect the performance of the entire ecosystem. Moreover, it depends on the maturity of the DW vendor to utilize that mechanism efficiently.

Disk-Memory-CPU Dance

The IT industry has used the term "in-memory" since the beginning of computer architectures. Ever since, differences of processing data in main memory versus storing data on external storage have dictated the limits on speed and capacity. The data to be processed usually exceeds the capacity of memory, which is always the case with big data.

Let us consider a simple analogy to understand the scale of this problem.

Imagine that you want a glass of milk to go with your sandwich. You go to the refrigerator, retrieve the milk, and pour a glass. It takes roughly 35 seconds. Next week, you go to the refrigerator, and there is no milk. You really want that glass of milk, so you go to the local grocery store to buy milk. It takes, say, 35 minutes (or 60 times longer) for that glass of milk. You are hungry, but the cold milk tastes so good!

If data is like the milk, the ratio of accessing data from memory versus disk is extreme—10,000 times! This ratio implies that it will require four days to get your glass of milk since the nearest grocery store is in Chicago while you reside in New York City. It totally changes the problem. Be glad that computers possess enormous patience!

Hence, a complex dance is required to move data from disk to memory, process calculations, and move results from memory to disk. Since big data has recently become BIG data, the dance can often last for a long duration, which is unacceptable to the business. Buying more memory does help, but memory is much more expensive than disk. Moreover, not all data has equal importance in its urgency for processing so that a dynamic balancing of memory-disk usage occurs. For example, the Teradata Intelligent Memory™ feature uses most-frequently-used caching algorithms to maximize the probability of having the proper data in memory when it is required.

In-Memory Optimization

This brings us to in-memory optimization (IMO), which focuses on the execution pipeline once the right data is in main memory. The analogy is like putting a milk dispenser on the refrigerator door so that there is an optimized path for accessing the milk.

Current processor architectures (like Intel Haswell AVX processors) have several high-speed memory caches between the CPU and main memory (as shown in the figure). With proper formatting, more data can be contained in these high-speed caches, reducing data movement from main memory. In addition, the architecture supports wider registers that hold more data so that the vector instructions can process a set of data in one CPU operation.

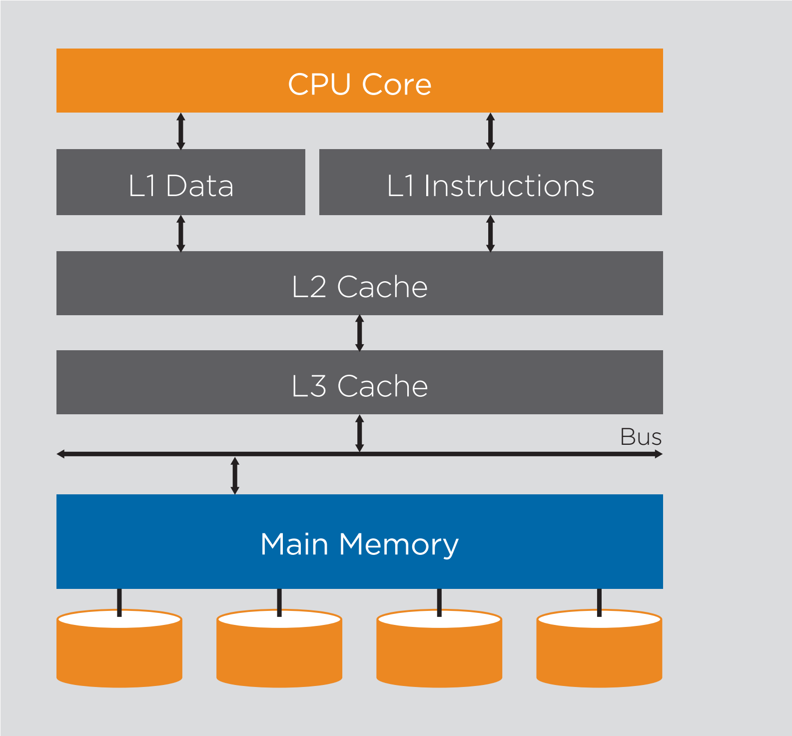

These new architectures provide the capability to process database operations with much greater efficiency. However, database system vendors must reengineer their low-level memory management to take advantage of these capabilities.

Teradata Database 15.10, in particular, supports a variety of in-memory optimization features, most of which require no changes by customers.

An example is to structure data in main memory into columnar formats (i.e., vector of values from a column) so that higher-level database operations can leverage vector instructions. Thus, a single CPU instruction cycle can process batches of data through bulk row qualification or hash joins, instead of a single row at a time. In other words, a single cycle can qualify batches of multiple rows for a table scan or link batches of join paths among tables. The database designer can optionally enhance this efficiency by defining certain columns as 'columnar' to better stream columnar vectors into main memory.

Another example is the capability of pipelining intermediate query results between processes by retaining data in main memory between process executions. Thus, the results of one database operation are passed directly to subsequent operations, avoiding storing intermediate results in spool on disk.

Do these low-level hardware advances as leveraged by Teradata's IMO features make a significant difference in performance of a large-scale database system? The next section documents a recent customer experience that observes these differences.

Customer Experience

A large global manufacturer employs over a hundred thousand persons, generating over $50B in annual revenue. The company has been a Teradata customer for over ten years, operating an enterprise-class data warehouse managing over 300 TB and supporting over 10 thousand daily users across all functional groups. The production platform centers on a 32+16 node Teradata Active Enterprise Data Warehouse 6700 with 512 GB of node memory, along with several other systems for development and testing.

> "Just about everyone" relies on the data warehouse.

Both business and IT users rely heavily on the data warehouse, via an internally developed query application, plus several BI tools like Tableau® and SAP® BusinessObjects. The company also utilizes Teradata Data Lab, which the company considers "a huge success" for prototyping new applications.

The company conducted a benchmark test to assess whether to upgrade their production Teradata Database system to version 15.10.

Benchmark Testing

The benchmark was part of their standard testing procedure to upgrade their production data warehouse to the Teradata Database 15.10. In addition to new features and bug fixes, this benchmark focused on the performance improvements from the in-memory optimizations (IMO) features in version 15.10.

The Teradata Database test system, consisting of 2+1 nodes, was used to compare their standard test workload on version 15.00 with the same workload on version 15.10. The DBA involved noted "Same hardware. Same settings. No SQL tuning. Just enable IMO. Refreshed DB stats, without adding any."

The benchmark consisted of a broad mix of ETL and BI queries from tactical to highly complex and large. The workload consisted of 512 queries, 800 tables, and 16TB of data, which was executed as two concurrent sessions. The company considered this benchmark as representing "real production workload as closely as is possible".

The company concluded that, when comparing results of the two workloads, the only contributing factor to performance was the IMO features. Since the benchmark did not include columnar-structured data tables, the columnar-related IMO feature was not tested.

Benchmark Results

As verified by the company, the results indicated an overall reduction in total runtime, I/O count, and spool usage by 29%, 35%, and 33% respectively. Individual query stats varied. For example, 7% (34 out of 512) showed greater than 10x improvement in CPU time, while 25% showing a small increase.

Once we reviewed details for the benchmark results, discussion shifted to the business implications. The issue was… So what? How does a 30% performance improvement impact the business?

Benchmark

- Teradata Database 15.00 vs 15.10
- Focus on in-memory optimizations
- 512 queries on 800 tables in 16TB

Results

- Runtime reduced by 29%
- I/O count reduced by 35%
- Spool usage reduced by 33%
- 7% of queries executed >10x faster

Improving User Experience

It was soon apparent that the primary goal for the company was improving user experience within an intensely utilized enterprise data warehouse. When probed, the team defined this experience as consisting of both responsiveness and consistency.

> "We are excited about these optimizations since it provides speed-of-thought responsiveness."

They did not define responsiveness in terms of seconds but as "speed-of-thought," which implied a smoothness (i.e., lack of interruption) in performing daily tasks. Because the EDW was mission-critical to the daily business operations, speed-of-thought responsiveness was a benefit visible through all levels of the organization.

Managing Peak Periods

When probed about whether a response time of two seconds versus three seconds was significant, a lively discussion occurred about peak usage windows from the internally generated workload. Because most of corporate functions are located in a single time zone, employees heavily use the system during the same hours. Thus, a 30% performance improvement would affect those peak windows positively.

> "We are getting breathing room to handle peak periods."

They summarized the benefits as, "getting breathing room to handle peak periods" to provide better user experience. Specifically, consistency of user experience was described as, "If it takes 20 seconds every day, but Monday morning it takes 5 minutes, then users are upset." Hence, managing EDW performance during those peak periods is critical to user experience.

Leveraging Better Value

The DBA team was initially motivated to upgrade to version 15.10 for new features and bug fixes. Hence, the performance improvements from IMO features were viewed as "gravy for us". They noted that there was no additional cost or effort for the company to "turn on" the IMO performance improvements. This successful benchmark test provided sufficient justification to use the IMO features in production.

> The performance improvement from IMO was "gravy for us."

When asked about their prior expectations of the results, they were "expecting a lot" since colleagues at two European companies reported similar performance improvements when upgrading to version 15.10.

Assessing the Benchmark

From the perspective of an independent industry analyst, the interview revealed a professionally conducted benchmark showing performance improvements from the IMO features. In an open discussion, two technical managers clearly explained the benchmark results, along with its issues. We debated the low-level technical details about MPP query processing, along with high-level business implications of user satisfaction.

One issue that was explored thoroughly is the size differential of test versus production environments. The test system had only 2+1 nodes while the production system has 32+16 nodes. They admitted, "The test workload is isolated and light, while the production box always has something running with a whole lot of contention. It may be too early to tell how these results will translate into actual production EDW with high concurrent workloads." However, the conclusion was that, since the company had successfully performed prior system upgrades using the same testing scheme, the size differential is not a concern.

Conclusions

In-memory technology has become a permanent component for large data warehouse systems. The progressive vendors, like Teradata Corporation, will continue to evolve this technology to provide greater capabilities. Hence, in-memory optimization is yet one more issue that IT managers will need to understand and utilize properly.

The company performed the benchmark adequately and assured itself that the upgrade to production status would proceed smoothly. They concluded that user experience using their EDW would have better speed-of-thought responsiveness and peak-period consistency, thus having positive impacts on employee satisfaction. They considered the IMO performance improvements "as gravy" since there was no cost or risk from changing their production workload. Further, the company predicts that their same EDW configuration (with minor hardware upgrades and database tuning) will support business operations into the future.

> "These performance gains will greatly improve business management confidence in the architecture and practice of our unified Enterprise Data Warehouse on Teradata."

"These performance gains will greatly improve business management confidence in the architecture and practice of our unified Enterprise Data Warehouse on Teradata."

Based on the experience of this company, IMO has business implications justifying the attention of business executives. In particular, the cost and effort of incorporating IMO into production EDW operations will have significant performance improvements and, thus, a clear business value impacts in several areas, such as employee satisfaction and system sustainability.

So for business executives, the next time your IT director discusses bulk row hash joins, listen and probe the business impacts. And for IT directors, start with the expected business impacts and explain how IMO technology will enable them.

About the Methodology

The methodology for this study is to listen carefully to pioneering companies in enterprise data warehousing and business analytics and to report accurately on their experience and perceptions. The intent is to contribute to professional education—to share the experiences and best practices with other IT professionals so that industry practices can mature, amid escalating business challenges and rapidly evolving technology.

In this study, a Teradata customer shared details of benchmark comparison of Teradata's latest in-memory optimizations, under conditions of anonymity. During a one-hour face-to-face interview, we discussed the objectives, results, and implications of this comparison. This study is an objective assessment of that comparison, based on that interview.

The primary author is Dr. Richard Hackathorn of Bolder Technology, who has been an independent industry analyst for several decades. Imad Birouty and Youko Watari of Teradata deserve credit for their substantive contributions and constructive criticism to this study.

About Bolder Technology

Bolder Technology Inc. is a twenty-year-old consultancy focused on Business Intelligence and Data Warehousing. The founder and president is Dr. Richard Hackathorn, who has more than thirty years of experience in the Information Technology industry as a well-known industry analyst, technology innovator, and international educator. He has pioneered many innovations in database management, decision support, client-server computing, database connectivity, and data warehousing.

Richard was a member of Codd & Date Associates and Database Associates, early pioneers in relational database management systems. In 1982, he founded MicroDecisionware Inc. (MDI), one of the first vendors of database connectivity products, growing the company to 180 employees. Sybase, now part of SAP, acquired MDI in 1994. He is a member of the Boulder BI Brain Trust (BBBT). He has written three books and has taught at the Wharton School and the University of Colorado. He received his degrees from the California Institute of Technology and the University of California, Irvine.

About Our Sponsors

Teradata (NYSE: TDC) helps companies get more value from data than any other company. Teradata's leading portfolio of big data analytic solutions, integrated marketing applications, and services can help organizations gain a sustainable competitive advantage with data. Visit **Teradata.com**.

03.16 EB9292

Bolder Technology, Inc.