



TERADATA INTELLIGENT MEMORY

TERADATA

Maximizing
Performance While
Increasing the
Value of Memory

TABLE OF CONTENTS

- 2 Executive Overview
- 3 Hottest Data Available for Maximum Performance
- 3 Teradata Intelligent Memory
- 4 Faster Queries, Better Decisions
- 4 Intelligent, Automatic, Transparent, and Efficient
- 6 System-wide Multi-Temperature Data Management
- 7 Conclusion

EXECUTIVE OVERVIEW

Teradata® Intelligent Memory speeds data warehouse query performance and maximizes the value of system memory by ensuring that the most frequently used data is kept in memory. It is Teradata’s latest innovation in its comprehensive, system-level approach to multi-temperature data management.

Available on all members of the Teradata workload specific platform family, Teradata Intelligent Memory is the only memory technology that automatically and transparently puts the hottest, most frequently used data in memory. This new data warehouse software was introduced with Teradata Database 14.10.

Teradata Intelligent Memory:

- ~ Provides a new extended memory space
- ~ Improves query performance
- ~ Is a smarter approach than in-memory databases
- ~ Leverages large memory capacities in new platforms

Faster Transactions with Intelligent Memory – In a March 2013 blog on big data and in-memory technology, Dr. Robin Bloor of The Bloor Group concluded the following:

“The ‘data analyst transaction’ is not a simple one. It varies significantly according to the goal and the nature of the data being analyzed. We cannot simply model it in the way we can model an OLTP transaction. But we know a couple things for sure. First, the transaction will go faster if the most frequently accessed data is held in memory and only has to be read from disk once. Second, it will go faster if we employ as much parallelism as possible.

And this means that in-memory technology and big data, whether they like it or not, must play nicely together.”

HOTTEST DATA AVAILABLE FOR MAXIMUM PERFORMANCE

As companies simultaneously expand near-real time tactical use of data warehouses, and store more data from traditional and new big data sources, faster query performance is imperative. But with the volume, complexity, variety, and variability of data—as well as uses—growing exponentially, delivering rapid query performance can be challenging.

In-memory databases try to solve the performance issue by storing all data in memory. However, this drives the cost of the system up and limits the amount of data that can be stored and made available to the business.

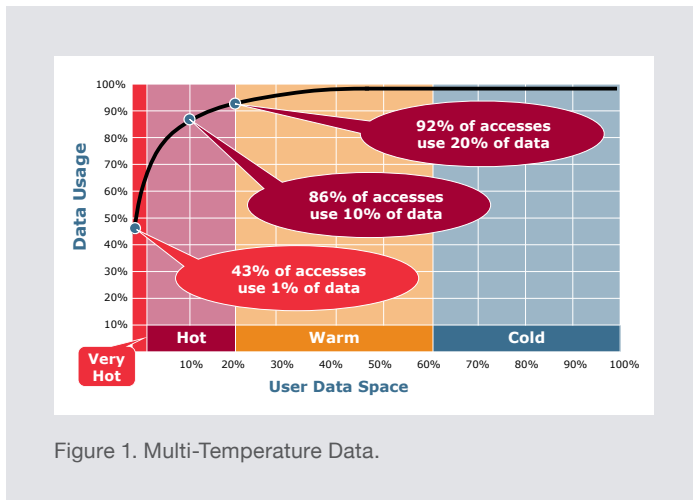


Figure 1. Multi-Temperature Data.

Teradata solved the problem by recognizing that all data does not have the same value and is not used in the same way. Rather than blindly storing all data in memory in a brute force attempt to meet data warehouse performance requirements, Teradata extended multi-temperature data management into the memory layer. Close to 50 percent of the query activity accesses just 1 percent of the data in many data warehouses. Teradata observed this relationship studying actual

system activity in a large cross-section of data warehouse customers. Teradata Labs leveraged this real-world behavior by designing Teradata Intelligent Memory to keep the hottest data in memory to achieve in-memory database performance without their cost.

TERADATA INTELLIGENT MEMORY

Teradata Database continuously tracks the temperature (or the relative access frequency) of all data in the database. The most frequently used (or hottest) data is identified internally within the database on a “very hot” data list. This list is modified and kept current as data warehouse workload and query access patterns change.

Whenever data on the very hot data list is read from disk during query processing, a copy is kept in Teradata Intelligent Memory’s extended memory area for future use. When data in Intelligent Memory is needed by another query, Teradata Database automatically looks to Intelligent Memory, eliminating the need for solid state disk (SSD) or hard disk drive (HDD) I/O, which is up to 3,000 times slower than memory access. If hotter data is used, it replaces the coolest data in Intelligent Memory so that the very hottest data is always available in memory for fast query performance with the lowest possible amount of disk I/O. But, since data is often used actively for an extended period of time, data is frequently kept in Intelligent Memory for hours, days, or even weeks at a time and used many thousands of times without further physical disk I/O.

With typical data warehouse query activity, a high percentage of queries can be using the data placed in memory by Teradata Intelligent Memory. However, since all data is still kept on disk storage and available for use, the data warehouse continues to meet the depth and breadth of its potential. Restrictions are not placed on user queries or data capture and retention in order to meet an artificial limit of storing all data in memory. Users get the answers they need with the performance they want because Intelligent Memory reduces the I/O needed in the system and increases performance and throughput.

FASTER QUERIES, BETTER DECISIONS

With the very hottest data in memory speeding query performance, users get better and more timely insight, make better decisions, and the organization gets more value from its data and data warehouse investment. The value of data warehouse applications increases and new applications are deployed with dramatic query performance improvements.

For instance, current sales data is used for customer service, inventory and stocking decisions, and even staffing projections. This generally keeps data for the current sales period very hot, improving the performance of all of these applications dramatically, because the database accesses the needed data from memory. The current period is also compared with the same period a year ago for trend and organizational performance analysis against goals. For this reason, the “year ago” reference period data may also be very hot and copied to memory by Teradata Intelligent Memory.

The sales data from 11 months ago is not the primary subject of management reports or current operational decisions and therefore likely won't be among the hottest data in the system. However, when an exception such as a warranty service request occurs and that data is needed, it is available on economical disk storage. In a month, when that same data becomes the reference data for many management reports and is accessed more frequently, it may become hot enough to be on the very hot list and be automatically copied into memory by Teradata Intelligent Memory.

With better query and application performance, and dynamic response to changes in query behavior and access patterns, application developers and users are free to use data in more and creative ways. It may become feasible to do more offer-customization to increase conversion rates on the website—or a warranty and quality model may be run more frequently to better adjust prices and increase margins once each model run takes less time.

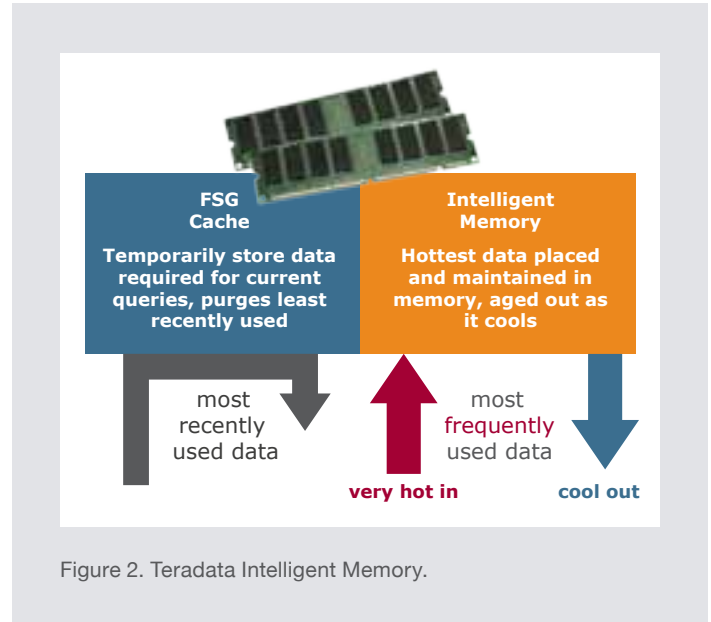


Figure 2. Teradata Intelligent Memory.

INTELLIGENT, AUTOMATIC, TRANSPARENT, AND EFFICIENT

Teradata Intelligent Memory goes about its business efficiently without the manual help of Database Administrators (DBAs) and transparently so that applications and users benefit from increased performance without changes to queries, applications, tools, or new training.

- ~ **Automatic intelligence:** Teradata Database continuously tracks the temperature of all data in the database automatically in the background. Teradata Intelligent Memory automatically copies data to its extended memory area whenever data on the very hot list is retrieved from disk for query processing. No DBA intervention is required to identify very hot data or to tell the system which data should be copied to memory and which data should only be kept in its persistent storage location. Teradata Database file system knows what data is available in memory and automatically uses that copy just as it would use data out of cache instead of going to disk. So no query changes are necessary and no special query syntax is used to take advantage of Teradata Intelligent Memory. If the needed data is in memory, that copy will be used.

“Teradata’s new in-memory architecture is integrated with its management of data temperature. This is very significant, because the hottest data will migrate automatically to the in-memory layer—Teradata Intelligent Memory; the next hottest data will move automatically to solid state disk; and, so on. Teradata also provides the column storage and data compression that amplify the value of data in memory. The customer sees increased performance without having to make decisions about which data is placed in memory.”

– Richard Winter, Chief Executive Officer, WinterCorp

- ~ **Transparent:** Because Teradata Intelligent Memory is transparent, no changes are required to query syntax, applications, business intelligence (BI) tool choice, user training, or physical database design. But some other physical database design performance techniques may no longer be necessary once Teradata Intelligent Memory is providing increased performance.
- ~ **Granular:** Teradata Intelligent Memory operates at the data block level, as opposed to a table or table partition level. Data temperature is tracked for small units of storage of about 11MB, called cylinders, within the Teradata file system for efficient granularity. Data blocks up to 1MB from the cylinders on the very hot list are copied into memory as they are accessed. This provides fast access to just the heavily used data without all of the cooler data from massive tables taking up expensive memory space. It also makes Teradata Intelligent Memory responsive to access pattern changes on even small areas of data due to business cycles and workload changes. Since temperature is tracked for a relatively small slice of data, a true and current temperature picture is maintained, rather than an average temperature over a range of temperatures in the very large amount of data in a table or partition.
- ~ **Multiplied value:** If more data can be stored in the Teradata Intelligent Memory extended memory area, more queries will be satisfied from memory, which is thus leveraged for even more value. For several years, Teradata Database has included several data

compression technologies to reduce the size of data when stored on disk. When compressed data are copied into Intelligent Memory, they are kept in their compressed format. Therefore, even more data can fit into the available memory.

- ~ **Columns or rows:** Teradata Columnar stores the data for a logical table column together on disk, increasing potential compression and making the query activity for the data stored in a block more homogeneous. These benefits of Teradata Columnar also automatically transfer to Teradata Intelligent Memory, further enhancing the granular and dynamic nature of Teradata Intelligent Memory and getting even more value from the available memory. If a table or part of a table is columnar on disk, that same data, organization, and highly compressed format will be copied into memory. It will be available for efficient query processing just as if it was read from the columnar table on disk, but without the I/O. In addition, since the data from a single column will dominate or fill a cylinder, the temperature measurement will be based on the use of that column for a very granular and responsive temperature to end user uses of the data. For example, if the customer number and product code in a sales table are used extensively, but the sale price is used infrequently, the temperature for product code data will be higher than the temperature for sale price in a columnar sales table. This leads to a precise match between business query activity and data kept in memory automatically, leveraging and multiplying the value of memory.

- ~ **Cache partner:** Teradata Database FSG cache (i.e., FileSegment cache, the primary internal cache within Teradata Database) is an extremely efficient short-term database cache. Teradata File System keeps several classes of data likely to be needed in the immediate future available in FSG cache (e.g., temporary tables, spool intermediate results, indexes). This mechanism achieves high “cache hit rates.” As with most computer cache techniques, FSG cache is very fluid and data is kept there for very short periods of time often seconds or possibly minutes while related processing is being done. Teradata Intelligent Memory partners with FSG cache with a focus on long-term data usage. Its extended memory area works alongside existing FSG cache and provides a more stable, temperature-based collection of data which will satisfy many queries over an extended period of time. Teradata Intelligent Memory works with FSG cache to ensure that only one copy of data is kept in memory at any given time to efficiently use the available memory. When data is accessed from either FSG cache or Teradata Intelligent Memory, those accesses are automatically tracked along with disk I/O accesses to maintain accurate temperature measurements for all data in the system.
- ~ **Relief from I/O constraints:** Teradata Intelligent Memory improves performance by reducing the need for storage I/O to retrieve data. Storage I/O (read/write access of data from permanent storage) and CPU processing resources are the largest categories of system resources which determine system performance. When either of these resources is overused or unavailable, performance suffers. By keeping the hottest data in memory and eliminating the need for I/O when that data is used, query performance improves unless lack of CPU resources otherwise limits performance improvement. Reducing I/O requirements will have limited performance benefit if the current problem is really completely busy CPU. But, as a result of its low overhead design, Teradata Intelligent Memory shouldn’t hurt performance even on systems with high CPU utilization.
- ~ **Performance continuity and data availability:** Following the infrequent occurrence of a database restart, Teradata Intelligent Memory automatically copies data from the very hot list into extended memory once the database recovers and enables user log-ons. This makes the system available for normal use as quickly as possible, even before all of the hottest data can be copied back into Intelligent Memory, and

avoids the long outages experienced by in-memory databases that must wait until all data is loaded into memory before being available for all types of queries. On the other hand, since Intelligent Memory automatically fills its memory with very hot data as quickly as possible, performance continuity is quickly restored.

- ~ **Integrated:** Teradata Intelligent Memory is integrated within the Teradata Database and Teradata file system. True to the Teradata Database shared-nothing architecture, each unit of parallelism (Teradata Access Module Processor or AMP) has its own Teradata Intelligent Memory logic and memory space. Teradata Intelligent Memory works with and supports other features of the Teradata Database. All other features continue to operate as they did before regardless of whether data is accessed from memory, SSD, or HDD during query processing. Indexes continue to speed access to the right data, whether it is in Intelligent Memory or on disk. Data in Intelligent Memory is available for any use, including updates. As you would expect, Teradata Database’s multi-dimensional linear scalability is maintained, but at a higher level of performance. There is no need for a separate system or appliance to get in-memory performance.

SYSTEM-WIDE MULTI-TEMPERATURE DATA MANAGEMENT

Teradata Labs takes a system design approach to multi-temperature data management. Teradata Database continuously tracks access to all data in the system and maintains temperature measurements for all data system-wide. The temperature is used throughout the system for temperature-appropriate management of the data.

For instance, Teradata Virtual Storage migrates data to the most appropriate storage location based on data temperature. As the temperature of data changes, it automatically migrates the data in the background to a location more appropriate for its new temperature. Teradata hybrid storage consists of HDD and SDD within the same Active EDW member of the Teradata workload specific platform family. On systems with hybrid storage, hot data is stored on solid state drives (SDD) and cooler data is stored on relatively slower, but less expensive hard disk drives (HDD). Data migration is automatically

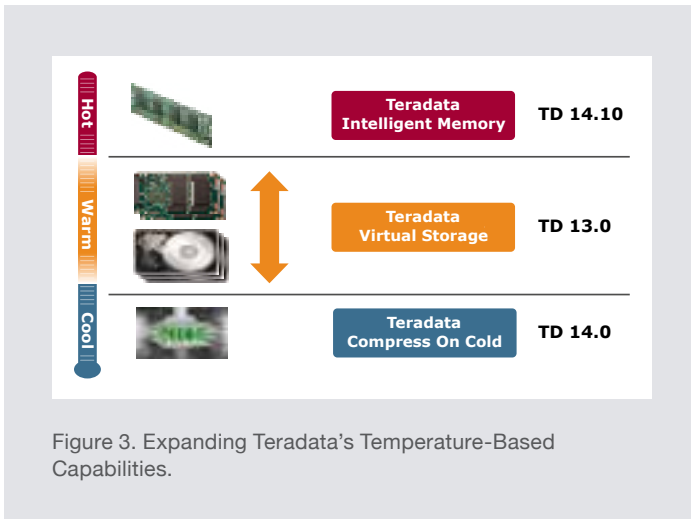


Figure 3. Expanding Teradata's Temperature-Based Capabilities.

managed as a background operation at levels designed to not interfere with database use. While the very responsive Teradata Intelligent Memory will copy the very hottest data to Intelligent Memory the next time it is accessed, Teradata Virtual Storage will migrate the permanent storage location of the very hot data and hot data to the relatively more abundant and economical SSD.

Teradata hybrid storage supports Teradata Virtual Storage by providing a wider range of storage characteristics to be taken advantage of by different temperature data.

Data temperature is also used by the Compress on Cold feature that automatically compresses cold data with Block Level Compression (BLC). BLC achieves very high compression rates, but consumes significant CPU processing resources compressing and decompressing the entire data block each time any data in the block is

stored or used. By automatically compressing cold data with BLC, significant disk space and I/O is saved, but CPU resources are rarely used because the data isn't accessed very often. If the data warms up, Compress on Cold automatically decompresses it and stores it in the decompressed form to avoid repeated CPU processing as long as the data is used frequently.

All of these multi-temperature data management features use the same temperature values as measured by the Teradata Database. Collectively, the features manage the data appropriately for its temperature, and implied usage, throughout the system.

CONCLUSION

With Teradata Intelligent Memory, Teradata continues to innovate and provide the highest performing Integrated Data Warehouse as part of the Teradata Unified Data Architecture™. Teradata Intelligent Memory uses large main memory to provide the highest query performance without the cost of in-memory databases. It delivers the best of both worlds: It keeps the hottest data in memory to achieve high performance—without the need to restrict available data to that which will fit in the available memory. With Teradata Intelligent Memory, Teradata Database continues to make the full scope of data available by keeping cooler data economically stored on disk.

In addition, Teradata Database delivers many other features and capabilities that provide high performance. These include the industry's best optimizer, efficient indexes, and several intelligent scan techniques to reduce the amount of data that must be read during query processing.

To learn more about Teradata Intelligent Memory and multi-temperature data management, visit Teradata.com.

