

# Teradata Warehouse 8.2 – speed, capability and ease

Enabling your enterprise to ask as many next questions as needed.

by Todd Walter

**“W**e never stop investigating. We are never satisfied that we know enough to get by. Every question we answer leads on to another question. This has become the greatest survival trick of our species.”

—Desmond Morris, British anthropologist, in *The Naked Ape*

I think this quote applies to organizations as much as to cultures or our species as a whole. Whenever an organization stops questioning, analyzing and learning, it begins the process of decline and death. An integrated data warehouse is a tool to support the never-ending quest to improve the organization in every dimension. Providing this tool to all stakeholders allows those next questions to be asked and answered. Teradata Warehouse 8.2 delivers new performance, capabilities and ease of management to allow the data warehouse to deploy more applications and answer more questions for more users than ever before.

## What is new in performance?

A new logging mechanism reduces I/O and improves throughput for integrating ever-increasing data volumes into the data warehouse. Updates that don't result in change are recognized. Index update and rollback performance enhancements reduce the cost of index maintenance. Join Indexes (JIs) can now be partitioned to significantly improve query performance.

## How is the logging mechanism changed?

The Transient Journal (TJ) is the Teradata mechanism for providing transaction consistency for structured query language (SQL) update operations. It captures the state of a row prior to an update operation so that it can be replaced (rolled back) if the update operation isn't completed and committed to the database. The TJ is enhanced with a new set of entries called the Write Ahead Log (WAL). These new entries allow update operations to be rolled forward as well as backward. Since all information necessary to apply the updates is captured in the log, the changed data blocks don't have to be written immediately. If a failure occurs before the blocks are written to the disks, the log is replayed to reapply the updates and bring the database to a consistent state.

The changed data blocks stay in the disk cache. Response time of the update operation may be reduced since it doesn't have to wait for the blocks to be written to disk before responding to the user. While there, the blocks may be updated many times before being physically written to the disk. Each of these additional updates saves an I/O operation, increasing the overall throughput of the system. The more frequent updates that occur to WAL rows in the same area of the data, the greater the system's performance; the biggest improvements will come from the most extreme volumes of SQL update operations.

## Doesn't "Buddy Back Up" eliminate the I/Os already?

The buddy mechanism made a copy of the changed data in another node in a Teradata massively parallel processing (MPP) system. This saved a disk I/O and protected the data in case of a node failure at the cost of a BYNET I/O and significant disk cache space in the buddy node. The buddy mechanism is removed completely in favor of the WAL logging, returning the disk cache space for use by other operations.

## How is the data kept safe from failures?

In certain complex failure scenarios, the buddy mechanism could result in complex recovery operations. All requirements to perform buddy flushes or FSGWizard recoveries are removed along with the buddy mechanism, reducing recovery times and increasing system availability. A "depot" is added to protect physical I/Os from partial failures while preserving the ability to update file system structures in place.

## Why do I care about updates that don't change the data?

We have found many cases where SQL update operations result in no change to the data in the database. Because it's easier for a tool or application to always perform the update than to track whether a change was actually made, the frequency of these no-change updates is increasing. Teradata now recognizes and optimizes these updates so they don't perform I/Os,

index changes or logging writes that would otherwise be required.

## How are index operations optimized?

Unique Secondary Index (USI) changes resulting from bulk SQL insert and delete operations change from row-by-row to block operations. As the rows in the table are changed, the changes to the index will be collected in a worktable. The worktable is redistributed by the rowhash of the USI, sorted and applied to each block of the USI in order. Many I/Os are eliminated, resulting in large performance improvements.

Non-Unique Secondary Index (NUSI) and USI are both logged in new ways to allow rollback to operate more efficiently. Many users avoid set processing updates or even drop indexes during update operations due to the risk of extended rollbacks. With the general rollback improvements in recent releases and this change for indexes, rollback operations are now expected to take the same length of time as the update operation. Update process strategies should be reconsidered and existing processes changed to take advantage of the removal of the risk of a long rollback when indexes are present.

## What is new for query performance?

The Partitioned Primary Index (PPI) feature is extended to JI to aid query performance. Partitioned Primary Index groups data together logically on disk, allowing narrow-range queries to read only data they need rather than the entire table. A JI optimizes access for high-frequency queries that utilize the data in a known way. Together they make a powerful tool for optimizing the high-frequency, narrow-range queries typical of front-line applications.

## What new capabilities does Teradata Warehouse 8.2 bring?

New data types, additions to triggers, IN list and identity column functionality allow new magnitudes of data and new applications to be created. Teradata enters a new era of openness with the

adoption of both Linux and Microsoft WS03 on 64-bit Intel EM64T platforms and replication is extended to work seamlessly among all Teradata operating system (OS) platforms.

## What new data types are available?

Enterprises are integrating ever more data and handling data from many countries. New data types—large-decimal and big-integer—allow the analysis of larger data sets and all the world's currencies. It's great (and I say that tongue-in-cheek) that very successful data warehouse customers require numbers up to 38 digits to account for their sales, products and customers.

## How are new applications enabled?

Business intelligence (BI) tools and applications have an endless appetite for larger pick-lists on more attributes. Internal limitations on optimization of combined long IN lists are lifted allowing this usage to benefit from all the query optimizations available to short lists. Real-time event applications utilizing triggers will no longer be restricted from the use of JI allowing the same tables to provide both high frequency query performance and event alerting to the users. Identity columns, traditionally used in batch extract, transform and load (ETL) processes, can now be easily used interactively with the addition of an interface that allows the new identity value to be returned when a new row is inserted.

## Tell me about Teradata on new OSs and 64-bit platforms

Intel's 64-bit EM64T technology combined with Linux or Microsoft WS03 to provide a new foundation for the Teradata solution. Teradata is ported to bring fully scalable, high performance, highly available and easily manageable data warehousing to these modern platform environments. For existing customers, the Teradata Database will continue to deliver all of its value on the MP-RAS SVR4 UNIX platform for the foreseeable future. All platform environ-

ments are supported by Teradata Replication Services allowing business continuity environments to be fully heterogeneous.

## What is new for those who have to manage the data warehouse?

CREATE TABLE AS is enhanced to copy statistics as well as data and ALTER TABLE can now change compression attributes on an existing table containing data, making table management easier. The database query log (DBQL) is enhanced to add statement type, utility objects, extensibility features and failed queries to enable detailed tracking of system operations. Spool space accounting becomes more accurate and correctly handles aborted requests, eliminating the requirement to monitor and clear these values.

Security is enhanced with a new password encryption algorithm and the ability to control logons by IP address. Workload management priority change is secured to only pre-defined priorities. System-level restarts are reduced by isolating faults in the dispatcher component to only the query that has an issue.

## Questions to ask next

Teradata data warehouses are all about asking the next question. Ease of definition, change and management allows new data to be analyzed quickly and easily. Data warehouse and analytic-centric features and functions allow applications to be built quickly and deployed widely. Highly scalable performance allows vast numbers of next questions to be asked by vast numbers of users with the latest information at their fingertips. Teradata Warehouse 8.2 extends the capability in each of these key areas allowing you to ask as many next questions as you need to help your enterprise thrive, not just survive. **T**

*Todd Walter, CTO, Teradata Development Division, oversees R&D efforts for Teradata Database software and systems. He is also responsible for the future vision and development of the active data warehouse. You can send your data warehousing questions to him at [todd.walter@teradata-ncr.com](mailto:todd.walter@teradata-ncr.com).*