

# COMPARING THE TERADATA SOLUTION AND NETEZZA

FIVE REASONS WHY NETEZZA DOESN'T MEASURE UP

TERADATA®

MARK SHAINMAN

Director of Competitive Programs, Teradata

## CASE STUDY: SUPERIOR SUPPORT ON A FLEXIBLE PLATFORM

For a large grocer supporting more than 1,600 stores in North America, the Netezza environment was a data duplication nightmare. The company had originally selected the platform to support strategic marketing functions, but the siloed Netezza environment could not integrate with a larger data warehouse infrastructure built on the Teradata platform.

The grocer chose a forklift migration to the Teradata solution. With a simplified data warehouse architecture based on an integrated infrastructure, the company was able to reduce the total cost of ownership of its data warehousing solution. Easy access to merged and integrated data empowers users to achieve far more detailed insight into business-critical data.

With attentive, affordable support from Teradata to help resolve even the biggest challenges, the grocer now manages more data with more confidence.

A Netezza data warehouse may seem like a great value—at first. Yet as data warehousing requirements grow over time, Netezza's performance tends to wane. The underlying architecture of the platform limits its ability to handle workloads beyond the sheltered confines of an analytical mart. The Teradata® solution, on the other hand, is designed to keep pace with the changing demands of today's big data environments.

## FIVE NETEZZA SHORTCOMINGS

Here are five key areas where Netezza simply cannot measure up to Teradata.

### 1. Concurrency

The latest Version 7 release of the Netezza database defaults to no more than 48 concurrent queries, and Netezza customers are actually encouraged to lower this number to improve response time. Netezza's hard limit of 63 concurrent transactions per system (including INSERT, UPDATE, DELETE, and data loads) can cause problems with extract, transformation, and load (ETL) activity. This limit forces administrators to keep a close watch on the number of concurrent ETL jobs. In contrast, Teradata Database can support up to millions of concurrent queries and transactions on a single system, allowing greater flexibility to meet complex demands.

Netezza has a single active host node that can be a performance bottleneck. All user sessions, all data loads, and all returned data must flow through this single server. Queries performing complex sorts and aggregations must send all data from potentially hundreds of back-end processes to this choke point for final merging and summarization. The Teradata massively parallel processing (MPP) architecture will automatically spread user sessions across multiple servers, providing scalable bandwidth for data flow. The Teradata solution can perform the final sorting and merging of results across the entire system using the patented intelligent BYNET software, eliminating the bottleneck of a single server.

### 2. Workload management

Although Netezza claims to be easy to use, its workload management scheme is surprisingly complex. The multilevel query throttling, multiple queues, and resource group settings can be a challenge to configure. This complexity contradicts Netezza's claims of simplicity, so much so that many Netezza customers do not go beyond the default workload management settings. IBM, which acquired Netezza in 2010, recommends that customers who run operational analytics with mixed workloads that need strong workload management should purchase IBM DB2 Universal Database, not Netezza.

In contrast, Teradata appliances offer robust workload management based on the same industry-leading technology as the high-end Teradata Active Enterprise Data Warehouse systems but with streamlined settings and controls. This allows Teradata appliance customers to benefit from the continuous improvement in workload management built into every new release of Teradata Database.

# COMPARING THE TERADATA SOLUTION AND NETEZZA

Many Netezza customers suffer from *box rage* because they are forced to maintain full copies of their data on two, three, or more Netezza systems.

## 3. Indexing

The Netezza approach to indexing is simple: It doesn't offer it. This is a major oversight because indexing is a key tool for performance on known or tactical workloads. Every data warehouse system has some planned reports that can be resolved with a well-chosen index. Instead, Netezza defaults to table scans and relies on its Zone Map feature for performance. This means that every new Netezza database release increases reliance on the physical data organization on disk for performance. For example, IBM claimed a twentyfold improvement in tactical query performance through use of the new page granular zone maps in Netezza 7.0, the latest release. However, this feature is only effective if data is sorted on disk and if you happen to be querying on the sort column. This means Netezza customers will be forced to frequently re-sort the data or maintain multiple copies of tables with different sort orders. This becomes complex to manage and adds a maintenance burden reminiscent of older database systems that continually need complex reorganization to stave off performance degradation.

In contrast, Teradata users have the option of creating indexes for common known queries. Teradata aggregate join indexes, a database-managed summary used to speed up queries, can dramatically improve performance for common ROLAP (relational online analytical processing) engines such as Microsoft Analysis Services or Oracle Essbase. Unlike other databases, the Teradata solution does not require the use of indexes, but it is good to have the option.

## 4. System expandability

Netezza customers cannot simply add nodes (servers) to an existing system to increase its size to meet growing capacity or performance needs. Instead, IBM must bring in an entirely new and larger Netezza system, copy the data over, and then remove the original system—in short, a floor sweep. IBM calls this complete system replacement an *expansion*. Netezza also does not allow customers to mix hardware generations in a single system. Limited concurrency and workload management give many Netezza customers *box rage* because they are forced to maintain full copies of their data on two, three, or more Netezza systems, with each serving a different application. The Teradata solution, by contrast, allows customers to expand systems incrementally by adding nodes (servers and storage) as necessary to meet growth needs. Teradata also offers investment protection by supporting the coexistence of multiple generations of hardware in a single system.

## 5. Leading-edge technology

Netezza's architecture inhibits its ability to leverage the latest advancements in hardware technology. Netezza's reliance on its proprietary field-programmable gate array (FPGA) accelerator technology prevents IBM from taking advantage of new technology. Even the latest N2001 (a.k.a. Striper) Netezza systems are still running the older 2011 vintage Westmere generation of Intel Xeon processor. Teradata was an early adopter of the Sandy Bridge generation and is now using the latest Ivy Bridge Xeon processors—a full two generations ahead of Netezza. Teradata adopted 2.5-inch disk drives in 2009 while Netezza waited until 2013. Teradata has been shipping high-performance solid-state drive (SSD) storage in data warehouse systems since 2009 and was the first vendor to ship an all-SSD appliance. In contrast, Netezza customers are still waiting for SSD technology. SAS chose Teradata as one of the first databases to support its High Performance Analytics MPP option. Although Netezza supports some SAS® in-database processing, it lacks the full range of SAS in-database capabilities such as SAS Analytics Accelerator and High Performance Analytics available on Teradata.

# COMPARING THE TERADATA SOLUTION AND NETEZZA

## TERADATA MIGRATIONS

The bottom line? In every area, from performance to functionality, Netezza falls short of the industry-leading capabilities available from Teradata solutions.

The good news is that Teradata has an excellent track record in managing migrations from Netezza. We follow proven best practices to make your transition as seamless as possible. And by moving to a superior data warehousing environment, you can cut operational costs, accelerate system performance, and gain access to a far broader range of mission-critical features.

So go ahead—make your move. Visit [www.Teradata.com](http://www.Teradata.com) or contact your Teradata account representative today.



10000 Innovation Drive Dayton, OH 45342 [teradata.com](http://teradata.com)

**TERADATA**<sup>®</sup>

Teradata and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or [Teradata.com](http://Teradata.com) for more information.