

At the core of the Teradata platform

New dual-core processor technology helps activate data warehousing. *by Jim Dietz*

Early in 2007, Teradata released the Teradata 5500 Server as its next generation of the Teradata massively parallel processing (MPP) platform. Pioneered almost 25 years ago, the Teradata platform is designed specifically to run the Teradata Database for data warehousing.

During the ongoing development of the Teradata Warehouse, Teradata saw unique platform system attributes were required to support the massive data and query workload of data warehousing applications. These five key attributes continue to drive the design criteria for this purpose-built platform family:

- > **Scalability.** Uniformly and efficiently increases capacity and performance to meet business needs
 - > **Availability.** Eliminates or dramatically reduces the impact of failures
 - > **Manageability.** Minimizes the human effort necessary to run the system
 - > **High-performance technology.** Rapidly exploits industry-leading technology
 - > **Growth with investment protection.** Expands and evolves the system while avoiding financial distress
- These purpose-built qualities become even more important as the enterprise data

warehouse (EDW) is extended to support not only the traditional knowledge workers in the back office but also operational front-line users in what Teradata calls active enterprise intelligence.

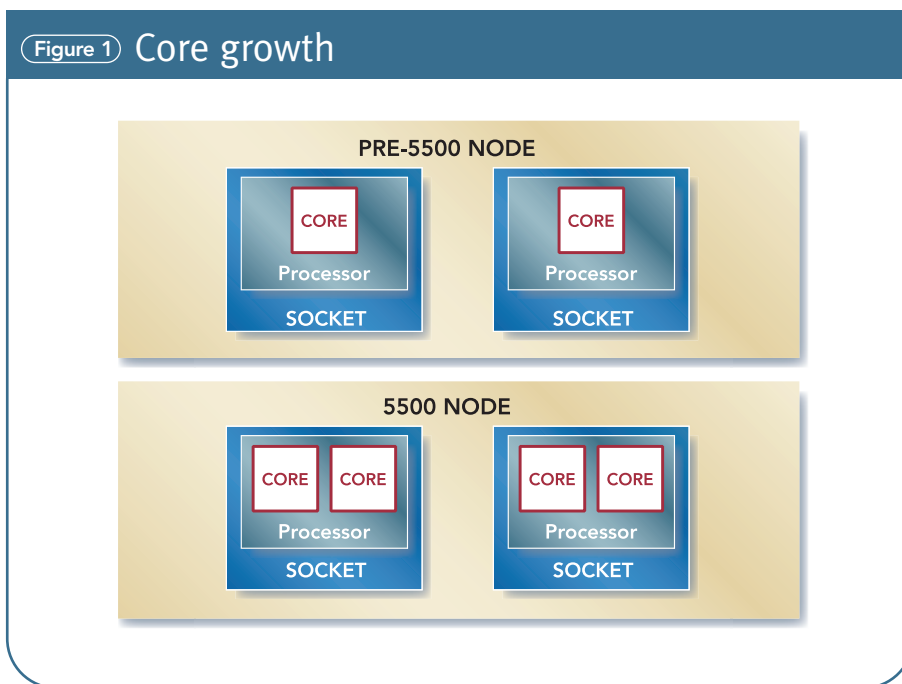
Faced with data tables that have exploded in size and users demanding near real-time decisions faster, the cycle-time compression for each of these active decisions magnifies the need for speed in the platform. The new platform technology delivered by the Teradata 5500 Server enables Teradata to meet the demands of this active data warehouse environment.

Core technology

The Teradata 5500 Server platform can be considered a field-proven product release because it is based on many of the prior generation Teradata 5450 Server components (such as BYNET interconnect, cabinetry and system management infrastructure). Like its predecessor, the Teradata 5500 Server can scale up in size to 1,024 nodes supporting up to 4 petabytes of data space.

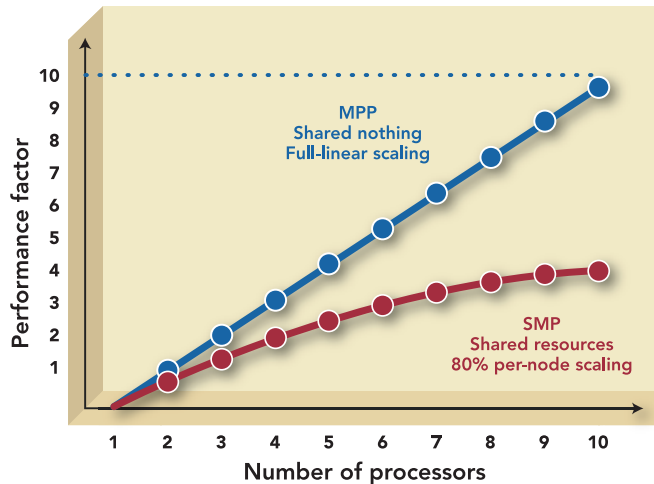
Before outlining the technology of the Teradata 5500 Server, two definitions are needed: “core,” which is a compute element that provides the basic functionality of a computing engine; and “node,” which refers to the individual servers in a Teradata platform system.

Prior to the Teradata 5500 Server, Intel Xeon processors used in previous generations contained a single core. With the advent of the Dual Core Intel Xeon Processor, now integrated into the



The Teradata 5500 Server has double the number of cores in each processor.

Figure 2 MPP and SMP performance differences



The massively parallel processing (MPP) data warehouse gains full benefit from each processor added, whereas the performance of the symmetric multi-processing (SMP) data warehouse stagnates.

Teradata 5500 Server platform, each processor contains two cores. Since Teradata servers can have one or two processors in each server node, the new Teradata 5500 Server node contains either two or four cores. (See figure 1, page 62.) With twice as many compute elements

(cores), the Teradata 5500 Server is capable of delivering a huge boost in node performance over previous generations.

On top of doubling the basic compute resources in a node, the Dual Core Intel Xeon Processor doubles the size of the cache memory on each processor in the Teradata

5500 Server and delivers a three- to four-fold increase in throughput speed for the memory and I/O buses that surround the processors in the node.

The 64-bit memory addressing provided by the Intel Xeon Extended Memory 64 Technology (EM64T) offers an additional boost in data warehousing power by extending beyond the 4GB memory limit of 32-bit technology. The additional memory (for instance, 8GB is recommended for the Teradata 5500 Server) enables a node to support more Teradata Database workload units, thereby delivering enhanced node performance. The benefits of the 64-bit addressing are available with either the Novell SUSE Linux or Microsoft Windows Server 64 operating systems, both of which are supported by the Teradata Database.

Teradata MPP: Maximizing system gains

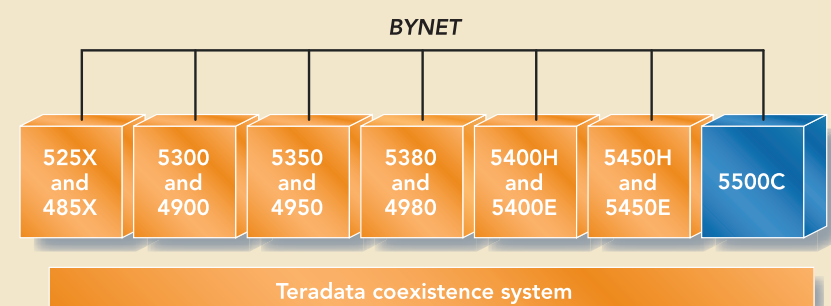
Beyond the gains with dual-core technology at the Teradata Server node level, Teradata has the unique ability to effectively apply nearly 100% of this node-level performance at the whole system level. The Teradata MPP architecture utilizes a true “shared

Coexistence protects the data warehouse investment

Teradata systems have the unique ability to grow incrementally over time using multiple generations of servers in the same system. This capability for a Teradata system to leverage the performance from each generation of servers in a mixed-generational system configuration is called “coexistence.” The Teradata 5500 Server, using the 5500C model, supports coexistence with up to six prior generations of nodes—including the Teradata 5250 Server, first released in 2000. With this mixed-generational capability, an investment in a Teradata Server today is protected for the future.

—J.D.

Seven generations in one system



Coexistence allows Teradata systems to build upon the performance of previous servers in a mixed-generational system configuration.

TPerf basics

The power of a data warehouse is its ability to do representative real-world data warehouse work, which Teradata measures in Traditional Performance (TPerf). Using a train analogy, the locomotive represents the power of the data warehouse, and the train cars represent the load (i.e., data). As depicted in the figure below, if a locomotive with a certain power level can pull two train cars, then pulling twice as many train cars at the same performance or speed will require twice the amount of measured power or locomotives. TPerf is the Teradata power metric to measure the power of the node, or locomotive, used to help configure the power needed to meet your data warehouse needs.

Teradata determines the TPerf of a Teradata Server by carefully measuring the performance of the three primary components (server node, memory and disk storage) while running the Teradata Database for basic data warehouse system operations. These operational measurements are combined into a representative enterprise data warehouse workload by measuring, weighting and summing each of the

operations that make up the workloads. These workloads include simple, complex and tactical queries, data loads and full table scans. The result is the TPerf metric for a server platform.

The other Teradata performance metric is related to speed. Using the same train analogy, a locomotive is capable of pulling a load of train cars at a set speed. If the number of cars is doubled, the same locomotive power will pull the train at half that speed. A train with half the number of cars allows the locomotive to pull at twice the speed.

Speed in a data warehouse relates to how much work the server node (i.e., the locomotive) can achieve per unit of an attached load, or data space. The Teradata Performance-per-capacity metric, called Perf-C, is simply the system or node TPerf divided by the attached user data capacity. Teradata Servers can be configured for various Perf-C performance levels based on the number and size of the disk drives attached to each node. This enables Teradata to consistently deliver industry-leading system performance that meets user needs.

—J.D.

nothing” approach to achieve full benefit for system performance as server nodes are added to the system. The data warehouse power of each Teradata Server node fully contributes to the total power of the Teradata system, regardless of the number of nodes in the system.

This near-linear scaling is enabled by the parallel architecture of the Teradata Database design and the distinctive BYNET interconnect. Other data warehouse solutions use combinations of shared resource approaches such as large symmetric multiprocessing (SMP) computers when applying dual-core technology. Consequently, their sharing of resources such as main memory limits the performance contribution of each server that is added to a system. Figure 2, page 63, illustrates the performance difference between the Teradata MPP and an SMP approach to system scaling.

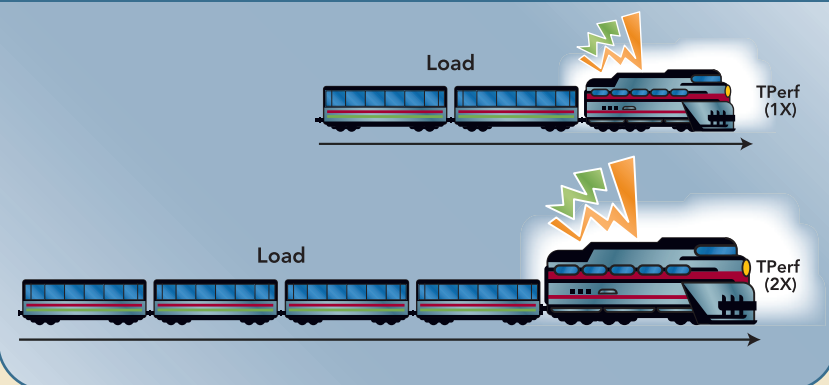
Predicting the power

A dependable method is needed to effectively configure and measure all of this data warehouse power now enabled with the Teradata 5500 Server. For nine generations of Teradata Server platforms, the Traditional Performance (TPerf) metric has been the basis for configuring and implementing hundreds of successful system expansions and new technology migrations.

Users of past generations of Teradata Servers have counted on the fact that they “pay for what they get” in the power of their data warehouse. The additional boost provided by the dual-core processor technology in the Teradata 5500 Server gives this relationship of price-for-power received even more impact.

To this end, the pricing of the Teradata solution software elements (database, tools and utilities) is based on the power delivered by the Teradata Server as defined by the server TPerf metric. This Teradata pricing approach unambiguously solves the industry conundrum on software pricing:

Measuring power through TPerf



A locomotive needs twice as much power to pull twice the load. Likewise, a data warehouse needs twice the power to perform at the same speed if the data is doubled. The Teradata Traditional Performance (TPerf) is a measure of the power in a Teradata Server.

how to fairly reflect the value of software performance on new processor technology.

Applying the power: A family affair

A broad family of Teradata 5500 Server products will facilitate the application of the dual core-enabled data warehouse power. Each of the three models leverages the key attributes of the purpose-built platform and is focused on a specific data warehouse need:

- > **5500H.** High-performance model that leverages two Dual Core Xeon processors and delivers the most power in the family
- > **5500C.** Coexistence model that uses one Dual Core Xeon processor to achieve compatibility with previous generations of Teradata servers
- > **5500E.** Entry-level model, with scalability limited to two nodes, that provides users a system with which to start small at a lower price and grow efficiently

The Teradata 5500 Server family models are compared in the table below. The user data space is the attached disk storage available for user-related data after accounting for overhead space. It is based on the recommended amount of storage for each node that delivers the most efficient balance of storage and node processing power.

'Green' power for activating the data warehouse

The combined performance and efficiency of the Teradata 5500 Server reduces electricity use by up to 70% compared with prior generations of Teradata nodes of just a few years ago with the same capability in user data and performance.

Additionally, the improved system density of the Teradata 5500 Server provides up to a 65% savings in data center real estate over previous generations. This resulting improvement in TPerf per watt and TPerf per floor space offers Teradata users a significant financial savings benefit in their data center.

With the introduction of the Dual Core Xeon Intel Processor, the Teradata 5500 Server is the first Teradata Server to leverage the compute performance of multiple-core processor technology. The data warehouse power provided through these processors meets the speed and time-compression demands of active enterprise intelligence enabled by an active data warehouse system. As a result, the Teradata 5500 Server has become the leading data warehouse platform for successful business decision making. **T**

Jim Dietz, Teradata platform marketing manager, has more than 12 years of experience with Teradata in developing and marketing data warehouse solutions.

Table: Purpose-built Teradata 5500 Server family

MODEL	5500H	5500C	5500E
Purpose-built attributes	• High performance • Scalability	• Investment protection • Scalability	Growth
Cores/node	4	2	2 and 4
System size	1 to 1,024 nodes	1 to 1,024 nodes	1 to 2 nodes
User data space/node	1.5TB	1TB	1TB or 1.5TB (2 core) (4 core)

A comparison of attributes and capabilities of the various Teradata 5500 Server models.