



WHITE PAPER

The Netezza Data Appliance Architecture:

A Platform for High Performance Data Warehousing and Analytics

Introduction

Success in any enterprise depends on having the best available information in time to make sound decisions. Anything less wastes opportunities, costs time and resources and can even put the organization at risk. But finding crucial information to guide the best possible actions can mean analyzing billions of data points and petabytes of data, whether to predict an outcome, identify a trend or chart the best course through a sea of ambiguity. Companies that can get this type of intelligence on demand are able to react faster and make better decisions than their competitors.

Continuing innovations in analytics can provide companies with an intelligence windfall that benefits all areas of the business. But when people need critical information urgently, the platform that delivers it should be the last thing on their minds. It should be as simple, reliable and immediate as a light switch, able to handle almost incomprehensible workloads without the complexity that gets in the way. It must also be built for longevity, with a technology foundation able to sustain performance as more users run increasingly complex workloads and data volumes continue to grow relentlessly, while offering the lowest total cost of ownership.

Extreme Performance with Appliance Simplicity

Netezza transforms the data warehouse and analytics landscape with a platform built to deliver extreme, industry-leading price-performance with appliance simplicity for years to come. It's a new frontier in advanced analytics, with the ability to carry out monumental processing challenges with blazing speed, without barriers or compromises. For users and their organizations, it means the best intelligence to all who need it — even as demands escalate from all directions.

The Netezza data warehouse and analytics appliance has a revolutionary design based on principles that have allowed Netezza to provide the best price-performance in the market. As a purpose-built appliance for high speed analytics, its power comes not from the most powerful and expensive components but from how the right components are assembled and work together to maximize performance. Massively parallel processing (MPP) streams combine multi-core CPUs with Netezza's unique FPGA Accelerated Streaming Technology (FAST™) engines to deliver performance that much more expensive systems cannot match or even approach. And as an easy-to-use appliance, the system delivers its phenomenal results out of the box, with no indexing or tuning required. Appliance simplicity extends to application development, enabling rapid innovation and the ability to bring high performance analytics to the widest range of users and processes.

This paper introduces Netezza's Asymmetric Massively Parallel Processin™ (AMPP™) architecture, and describes how the system orchestrates queries and analytics to achieve its unprecedented speed. We'll see how Netezza software and hardware come together to extract the maximum utilization from every critical component, and how a system optimized for tens of thousands of users querying huge data volumes really works. It's a unique data warehouse and analytics platform with unparalleled price-performance, ready for today's needs and tomorrow's challenges.

Architectural Principles

The Netezza appliance integrates database, processing and storage in a compact system optimized for analytical processing and designed for flexible growth. The system architecture is based on the following core tenets that have been a hallmark of Netezza's price-performance leadership in the industry:

Processing Close to the Data Source

The Netezza architecture is based on a fundamental computer science principle: when operating on large data sets, do not move data unless you absolutely have to. The Netezza fully exploits this principle by utilizing commodity components called Field Programmable Gate Arrays (or FPGAs) to filter out extraneous data as early in the data stream as possible, as fast as data can be streamed off the disk. This process of data elimination close to the data source removes IO bottlenecks and frees up downstream components such as the CPU, memory and network from processing superfluous data, thus having a significant multiplier effect on system performance.

Balanced, Massively Parallel Architecture

The Netezza architecture combines the best elements of SMP and MPP to create a purpose-built appliance for running blazing fast analytics on petabytes of data. Every component of the architecture, including the processor, FPGA, memory and network, is carefully selected and optimized to service data as fast as the physics of the disk allows, while minimizing cost and power consumption. The Netezza software orchestrates these components to operate concurrently on the data stream in a pipeline fashion, thus maximizing utilization and extracting the utmost throughput from each MPP node. In addition to raw performance, this balanced architecture delivers linear scalability to more than a thousand processing streams executing in parallel, while offering a very economical total cost of ownership.

Platform for Advanced Analytics

The principles of MPP and data processing close to the source are equally applicable to advanced analytics on large data sets. Netezza allows complex non-SQL algorithms to be easily embedded in the processing elements of its MPP streams without the typical intricacies of parallel or grid programming. The ability to run analytics of any complexity "on stream" against huge data volumes eliminates the delays and costs of moving data to separate hardware. It also accelerates performance by orders of magnitude, making Netezza the ideal platform for the convergence of data warehousing and advanced analytics.

Appliance Simplicity

The Netezza automates and streamlines day-to-day operations, shielding users from the underlying complexity of the platform. Simplicity rules whenever there is a design tradeoff with any other aspect of the appliance. Unlike other solutions, it just runs – handling demanding queries and mixed workloads with blistering speed, without the tuning required by other systems. Even normally time-consuming tasks such as installation, upgrades, and ensuring high-availability and business continuity are vastly simplified, saving precious time and resources.

Accelerated Innovation and Performance Improvements

One of the key goals of the Netezza architecture is to deliver price-performance improvements and innovative functionality faster than competing technologies over the long run. While the use of open, blade-based components allows Netezza to incorporate technology enhancements very quickly, the turbocharger effect of the FPGA, a balanced hardware configuration and tightly coupled intelligent software combine to deliver overall performance gains far greater than those of individual elements. In fact, Netezza has delivered more than 4X performance improvement every two years (double that of Moore's Law¹) since its introduction, far outpacing other well-established vendors.

Flexible Configurations and Extreme Scalability

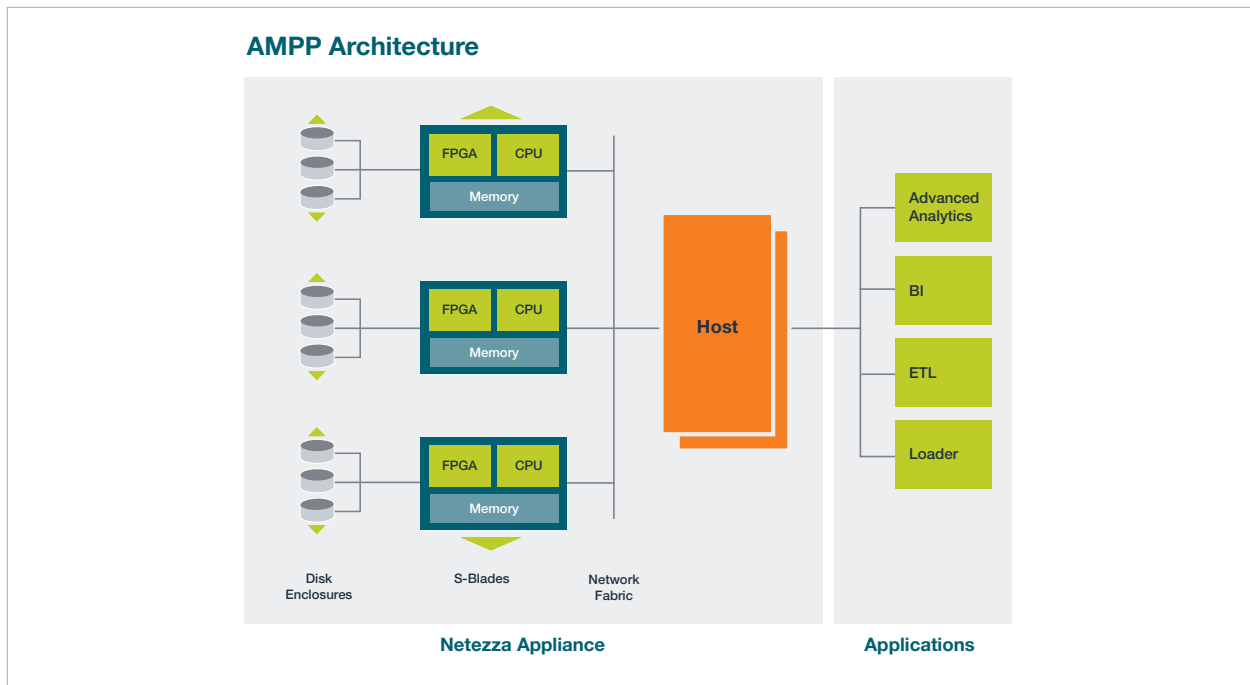
The Netezza scales modularly from a few hundred gigabytes to tens of petabytes of queryable user data. The system architecture is highly adaptable to serve the needs of different segments of the data warehouse and analytics market. The use of open blade-based components allows the disk-processor-memory ratio to be easily modified in configurations that cater to performance- or storage-centric requirements. The same architecture also supports memory-based systems that provide extremely fast, real-time analytics for mission-critical applications.

The following pages examine how Netezza puts these principles into practice.

1. Gordon Moore, Intel co-founder, predicted in 1965 that the number of transistors on a chip will double about every two years. Software applications generally rely on these processor improvements to accelerate performance over time.

System Building Blocks

A major part of Netezza's performance advantage comes from its unique AMPP architecture, which combines an SMP front-end with a shared-nothing MPP back-end for query processing. Each component of the architecture is carefully chosen and integrated to yield a balanced overall system. Every processing element operates on multiple data streams, filtering out extraneous data as early as possible. More than a thousand of these customized MPP streams work together to "divide and conquer" the workload.



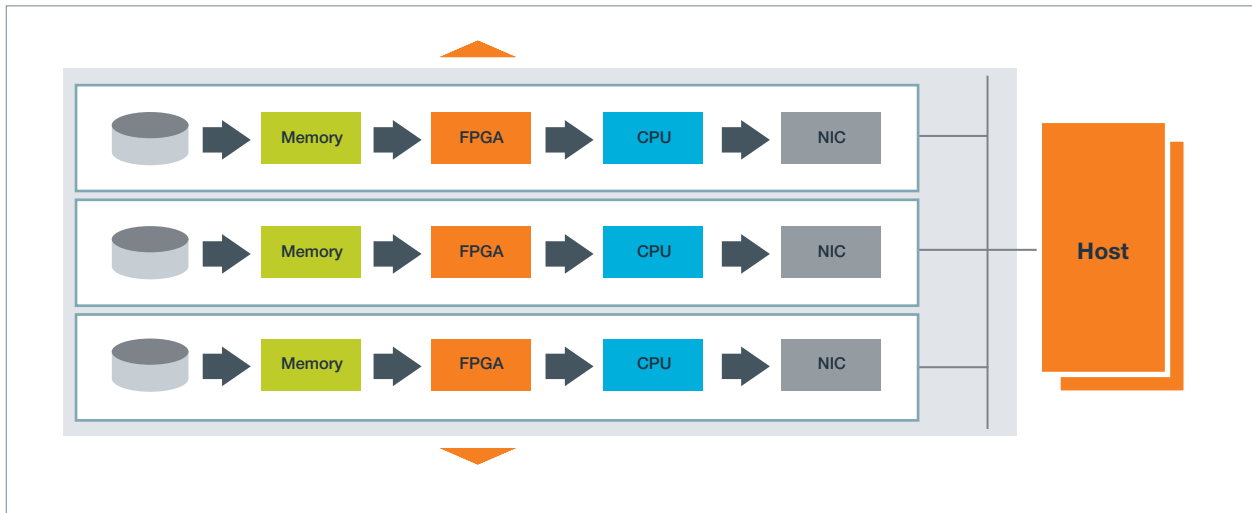
Let's examine the key building blocks of the appliance:

- **Netezza Hosts:** The SMP hosts are high-performance Linux servers that are set up in an active-passive configuration for high-availability. The active host presents a standardized interface to external tools and applications. It compiles SQL queries into executable code segments called snippets, creates optimized query plans and distributes the snippets to the MPP nodes for execution.
- **Snippet Blades (S-Blades):** S-Blades are intelligent processing nodes that make up the turbocharged MPP engine of the appliance. Each S-Blade is an independent server that contains powerful multi-core CPUs, multi-engine FPGAs and gigabytes of RAM, all balanced and working concurrently to deliver peak performance. The CPU cores are designed with ample headroom to run complex algorithms against large data volumes for advanced analytics applications.
- **Disk Enclosures:** The disk enclosures contain high-density, high-performance disks that are RAID protected. Each disk contains a slice of the data in a database table. The disk enclosures are connected to the S-Blades via high-speed interconnects that allow all the disks in a Netezza to simultaneously stream data to the S-Blades at the maximum rate possible.
- **Network Fabric:** All system components are connected via a high-speed network fabric. Netezza runs a customized IP-based protocol that fully utilizes the total cross-sectional bandwidth of the fabric and eliminates congestion even under sustained, bursty network traffic. The network is optimized to scale to more than a thousand nodes, while allowing each node to initiate large data transfers to every other node simultaneously.

Note: All system components are redundant. While the hosts are active-passive, all other components in the appliance are hot-swappable. User data is fully mirrored, enabling better than 99.99% availability.

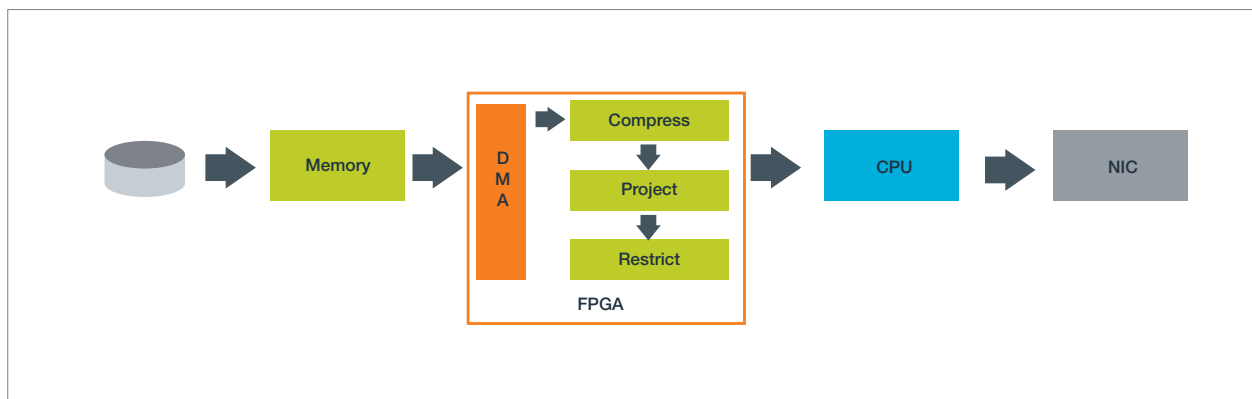
Where Extreme Performance Happens — Inside an S-Blade

A Snippet Processor (one of many): Commodity components and Netezza software combine to extract the utmost throughput from each MPP node. A dedicated high-speed interconnect from the storage array allows data to be delivered to memory as quickly as it can stream off the disk. Compressed data is cached in memory using a smart algorithm, which ensures that the most commonly accessed data is served right out of memory instead of requiring a disk access. FAST Engines running in parallel inside the FPGAs uncompress and filter out 95-98% of table data at physics speed, keeping only the data that is relevant to answer the query. The remaining data in the stream is processed concurrently by CPU cores, also running in parallel. The process is repeated on more than a thousand of these parallel Snippet Processors running in the Netezza. The result is performance that exceeds much more expensive systems by orders of magnitude.



Turbocharging the S-Blades: The Power of Netezza FAST Engines

The FPGA is a critical enabler of the price-performance advantages of the Netezza. Each FPGA contains embedded engines that perform filtering and transformation functions on the data stream. These FAST engines are dynamically reconfigurable, allowing them to be modified or extended through software. They are customized for *every snippet* through parameters provided during query execution and act on the data stream delivered by a Direct Memory Access (DMA) module at extremely high speed.



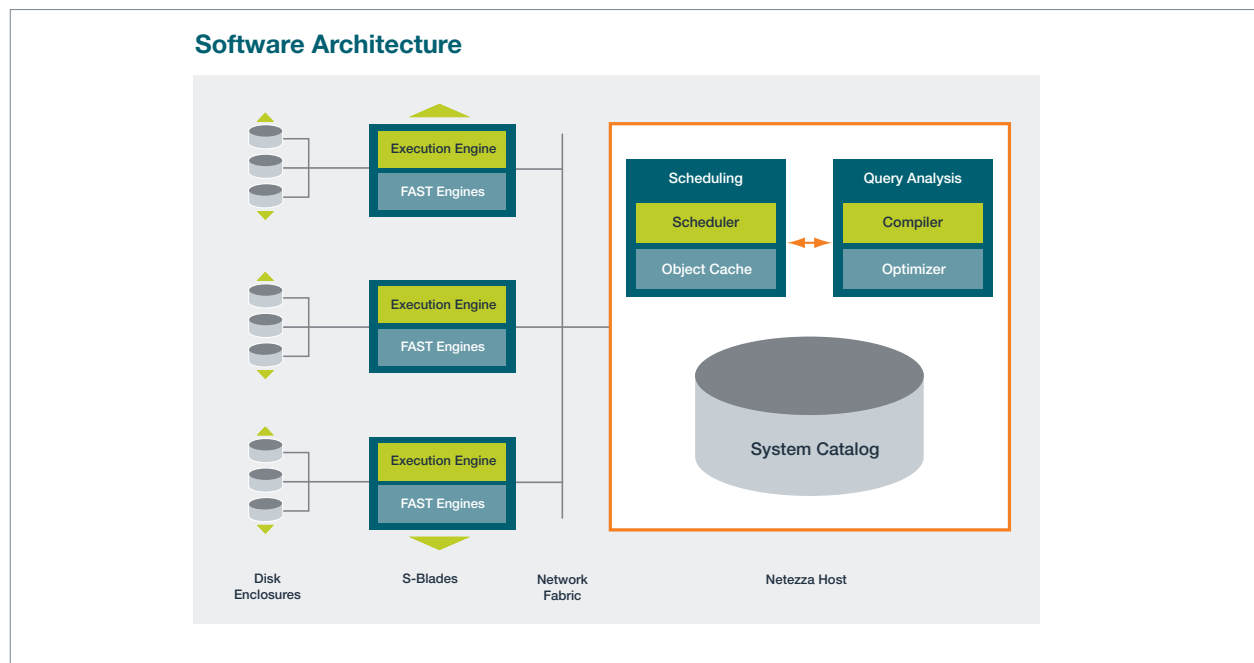
FAST engines include:

- The Compress engine, a Netezza innovation that boosts system performance by a factor of 4-8X. The engine uncompresses data at wire speed, instantly transforming each block on disk into 4-8 blocks in memory. The result is a significant speedup of the slowest component in any data warehouse — the disk.
- The Project and Restrict engines, which further enhance performance by filtering out columns and rows respectively, based on the parameters in the SELECT and WHERE clauses in a SQL query.
- The Visibility engine, which plays a critical role in maintaining ACID (Atomicity, Consistency, Isolation and Durability) compliance at streaming speeds in the Netezza. It filters out rows that should not be “seen” by a query; e.g. rows belonging to a transaction that is not committed yet.

The Netezza FAST engines also provide an extensible framework for innovative new functions to be added in the future through enhancements to the Netezza software. These new functions promise to boost system performance, security and reliability even further.

Orchestrating Queries on the Netezza

The Netezza’s hardware components and intelligent system software are closely intertwined. The software is designed to fully exploit the hardware capabilities of the appliance and incorporates numerous innovations to offer exponential performance gains, whether for simple inquiries, complex ad-hoc queries or deep analytics. In this section, we’ll examine the intelligence built into the system every step of the way.



Netezza software components include:

- A sophisticated parallel optimizer that transforms queries to run more efficiently and ensures that each component in every processing node is fully utilized
- An intelligent scheduler that keeps the system running at its peak throughput, regardless of workload
- Turbocharged Snippet Processors that efficiently execute multiple queries and complex analytics functions concurrently
- A smart network that makes moving large amounts of data through the Netezza a breeze

Let’s see how these elements work together, starting when a user submits a query. Technology-savvy readers will see that the Netezza processes queries very differently than other data warehouse systems.

Make an optimized query plan...

The host compiles the query and creates a query execution plan optimized for Netezza's AMPP architecture. The intelligence of the Netezza optimizer is one of the system's greatest strengths. The optimizer makes use of all the MPP nodes in the system to gather detailed, up-to-date statistics on every database table referenced in a query. A majority of these metrics are captured during query execution with very low overhead, yielding just-in-time statistics that are individualized per query. The appliance nature of the Netezza, with integrated components able to communicate with each other, allows the cost-based optimizer to more accurately measure disk, processing and network costs associated with an operation. By relying on accurate data rather than heuristics alone, the optimizer is able to generate query plans that utilize all components with extreme efficiency.

Intelligence in the Optimizer: Calculating Join Order

One example of optimizer intelligence is the ability to determine the best join order in a complex join. For example, when joining multiple small tables to a large fact table, the optimizer can choose to broadcast the small tables in their entirety to each S-Blades, while keeping the large table distributed across all Snippet Processors. The approach minimizes data movement while taking advantage of the AMPP architecture to parallelize the join.

The optimizer utilizes these statistics to transform queries before processing begins, in order to minimize disk I/O and data movement — the things that slow down performance in a data warehouse system. Transforming operations performed by the optimizer include:

- Determining correct join order
- Rewriting expressions
- Removing redundancy from SQL operations

Convert it to snippets...

The compiler converts the query plan into executable code segments, called snippets – query segments executed by Snippet Processors in parallel across all the data streams in the appliance. Each snippet has two elements: compiled code that is executed by individual CPU cores and a set of FPGA parameters that customize the FAST engines filtering for that particular snippet. This snippet-by-snippet customization allows the Netezza to provide, in effect, a hardware configuration optimized on the fly for individual queries

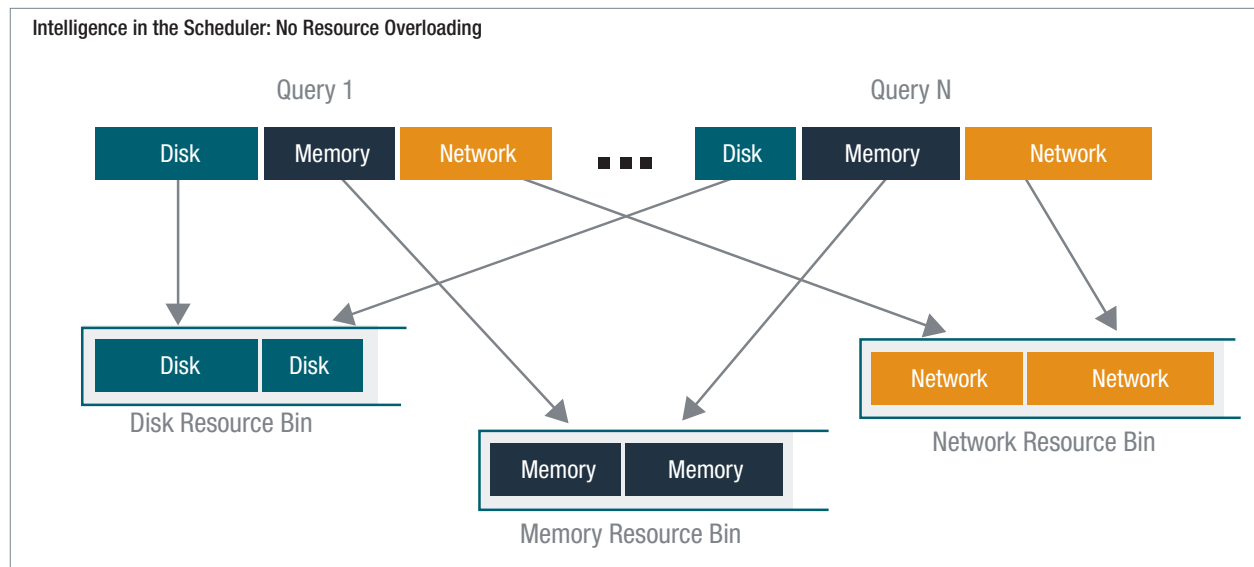
Intelligence in the Compiler: the Object Cache

The host uses a feature called the object cache to further accelerate query performance. This is a large cache of previously compiled snippet code that supports parameter variations. For example, a snippet with the clause, "where name='bob'" might use the same compiled code as a snippet with the clause, "where name='jim'" but with settings that reflect the different name. This approach eliminates the compilation step for over 99% of snippets.

Schedule them to run at just the right moment...

The Netezza scheduler balances execution across complex workloads to meet the objectives of different users, while maintaining maximum utilization and throughput. It considers a number of factors including query priority, size and resource availability in determining when to execute snippets on the S-Blades. The appliance architecture allows the scheduler to gather more up-to-date and accurate metrics about resource availability from each component of the system. The scheduler uses sophisticated algorithms that maximize system throughput by utilizing close to 100% of the disk bandwidth and ensuring that memory and network resources do not get overloaded and cause the system to thrash and lose efficiency. This is an important characteristic of the Netezza, ensuring the system keeps performing at peak throughput even under very heavy loads.

When the scheduler gives the green light, the snippet is broadcast to all Snippet Processors through the intelligent network fabric.



Execute them in parallel...

Each Snippet Processor on every S-Blade now has the instructions it needs to execute its portion of the snippet. In addition to the host scheduler, the Snippet Processors every have their own smart preemptive scheduler that allows snippets from multiple queries to execute simultaneously. The scheduler takes into account the priority of the query and the resources set aside for the user or group that issued it to decide when and for how long to schedule a particular snippet for execution. When that instant arrives, it's showtime:

1. The processor core on each Snippet Processor configures the FAST engines with parameters contained in the query snippet and sets up a data stream.
2. The Snippet Processor reads table data from the disk array into memory, utilizing a Netezza innovation called ZoneMap™ acceleration to reduce disk scans. The Snippet Processor also interrogates the cache before accessing the disk for a data block, avoiding a scan if the data is already in memory.
3. The FPGA then acts on the data stream. It first accelerates the data stream by a factor of up to 4-8X by uncompressing the data stream at wire speed.
4. The FAST engines then filter out any data that is not relevant to the query. The remaining data is streamed back to memory for concurrent processing by the CPU core. This data is typically a tiny fraction (2-5%) of the original stream, greatly reducing the execution time required by the processor core.

5. The processor core picks up the data stream and performs core database operations such as sorts, joins and aggregations on it. It also applies complex algorithms that are embedded in the Snippet Processor for advanced analytics processing
6. Results from each Snippet Processor are assembled in memory to produce a sub-result for the entire snippet. This process is repeated simultaneously across more than a thousand Snippet Processors, with hundreds or thousands of query snippets executing in parallel.

ZoneMap Acceleration – the Netezza Anti-index

ZoneMap acceleration exploits the natural ordering of rows in a data warehouse to accelerate performance by orders of magnitude. The technique avoids scanning rows with column values outside the start and end range of a query. For example, if a table contains two years of weekly records (~100 weeks) and a query is looking for data for only one week, ZoneMap acceleration can improve performance up to 100 times. Unlike indexes, ZoneMaps are automatically created and updated for each database table, without incurring any administrative overhead.

And return the results!

All Snippet Processors now have snippet results that need to be assembled. The Snippet Processors use the intelligent network fabric to communicate flexibly with the host and with each other to perform intermediate calculations and aggregations.

Intelligence in the Network:

Predictable performance and scalability

Netezza's custom network protocol is designed specifically for the data volumes and traffic patterns associated with high volume data warehousing. The Netezza protocol ensures maximum utilization of the network bandwidth without overloading it, allowing predictable performance close to the line rate.

Traffic flows smoothly in three distinct areas:


- From the host to the Snippet Processors (1 to 1000+) in broadcast mode
- From Snippet Processors to the host (1000+ to 1), with aggregation in the S-Blades and at the system rack level
- Between Snippet Processors (1000+ to 1000+), with data flowing freely on a massive scale for intermediate processing

The host assembles the intermediate results received from the Snippet Processors, compiles the final result set and returns it to the user's application. Meanwhile, other queries are streaming through the system at various stages of completion.

All Information on Demand, For Everyone Who Needs It

Sometimes the best solutions aren't the biggest or most expensive, but have the smartest design. Netezza recognized early on that streaming processing provides an inherent advantage over the traditional computing architectures used by other analytic and data warehousing systems. The result is a compact appliance with performance that dwarfs that of much larger systems, with blinding speed for running complex algorithms against huge data volumes and the mixed workloads of thousands of concurrent users. Processing performance is complemented by other capabilities that make the Netezza a unique platform to help businesses succeed:

- **Simplicity of use:** The Netezza is self-managed just like an appliance should be and always running at its peak throughput. The system software makes sure of that, without human intervention.
- **Better decisions across the enterprise:** Embedded functions bring a new generation of analytics into the database with minimum development effort. There's no need for separate server hardware or time lost in massive data transfers — just lightning-fast results and the ability to bring crucial business intelligence to everyone who could benefit, at all sectors of an organization.
- **Agility for the future:** The system is built not just for today's challenges, but for years to come, scaling linearly to tens of petabytes of user data and with performance acceleration far beyond the conventional speed-up governed by Moore's Law.

The Netezza allows users and their companies to make decisions with maximum clarity while taking performance for granted. But don't just take our word for it. The best way to appreciate the Netezza is to see it in action. We think you'll agree there's simply nothing else like it for making the most out of your data. 



SAFE HARBOR

"Safe Harbor" Statement under the U.S. Private Securities Litigation Reform Act of 1995: Certain information contained in this document is forward-looking in nature. Any expectations based on these forward-looking statements are subject to risks and uncertainties and other important factors. These and many other factors could cause delivery of products, features or enhancements to differ materially from expectations based on these forward-looking statements. Netezza does not undertake an obligation to update its forward-looking statements to reflect future events or circumstances.



Netezza Corporation

26 Forest Street, Marlborough, MA 01752

+1 508 382 8200 TEL

+1 508 382 8300 FAX

www.netezza.com

About Netezza

Netezza (NYSE: NZ) is the global leader in data warehouse and analytic appliances that dramatically simplify high-performance analytics across an extended enterprise. Netezza's technology enables organizations to process enormous amounts of captured data at exceptional speed, providing a significant competitive and operational advantage in today's data-intensive industries including digital media, energy, financial services, government, health and life sciences, retail and telecommunications. Netezza is headquartered in Marlborough, Massachusetts and has offices in Northern Virginia, Canada, the United Kingdom, Germany, France, Japan, Korea, Australia and Singapore.

For more information about Netezza, please visit www.netezza.com.