Company Brief – October 2002
Netezza

Huh?

Netezza is a well-funded startup focused on developing a purpose-built device for Business Intelligence applications. The crux of the issue is that it simply takes too long to extract real information from very large databases today. On the one hand, businesses are getting inundated with new data from web sites, CRM and other applications. On the other hand, businesses are required to make quick, fact-based decisions to be competitive. As these databases increase in size and complexity, traditional methodologies of using general-purpose systems, originally designed for OLTP applications, are failing to provide answers in a timeframe that is consistent with good decision-making. If Netezza has its way all this will change.

The Problem

In spite of the state of worldwide economy, the fact is the incoming data growth rates have continued to climb exponentially. The web sites are still capturing mounds of customer data (click streams are an incredible producer of raw data), satellites are capturing even more images, intensified oil exploration is producing terabytes of seismic data and rich media is playing havoc on storage systems. As we have pointed out in our Reference Data report, we are already in an era of data explosion. It will only get worse.

To harness this data, extract information from it and use this information in the decision making process, is to gain competitive advantage. To fail to do so is to wither away.

OK, so what's the big deal? Don't we have powerful systems and database software to deal with these? Yes, we do but there is an ever widening disparity between the rate at which data is piling up and the time it takes the most powerful existing systems to truly extract information in a timeframe that matters.

Let's say we are dealing with online retail. In order to recommend the most appropriate upgrades or additional products to an online customer, you want to be able to extract the buying pattern of this buyer from last two years of data in a 50TB database. If it takes the system even five minutes to determine the answer, you have lost the opportunity or at least reduced your chances of additional sales. An analysis done in five seconds would have proved significantly more valuable. But there are no affordable commercial systems today to do that.

Similarly, if it takes 250 hours of number crunching to determine the best actions to manage your inventories, your decision may be obsolete by the time the analysis is available. Results from an arbitrage analysis in 5 versus 50 minutes could mean millions of dollars in profits to a Wall Street company. There are hundreds of other examples in practically every industry.

The Current State-of-the-Art

The current solutions are based upon computer architectures that are now decades old. One can either buy the most powerful SMP (Symmetrical Multiprocessor) systems with massive amounts of main memory, cache and storage and run DBMS software on them. Or one could cluster several such systems with shared storage, or one could purchase specialized MPP (massively Parallel Processors) systems that break the problem down into many elements and let each be handled by a processor/memory/disk complement. NUMA (Non-uniform Memory Architecture) systems have been tried. Each one suffers from a different form of bottleneck. Some suffer from requiring programming

changes and constant fine-tuning. Companies desperate for results buy bigger and bigger systems in the hope of reducing the query time. But the payoff is incrementally smaller due to the inherent limitations of these architectures.

To be sure, in mid 80's there was an effort to build specialized database systems but these efforts failed, primarily because they added proprietary interfaces and required fine-tuning of the applications. Customers wanted faster results but they were not ready to give up the freedom they had begun to enjoy with open systems. So they relied on general-purpose systems of increasing performance but still got answers in less than "real time" (see definition later). Until now.

Enter Netezza

Netezza has developed a purpose-built system for speeding up SQL queries by one to two orders of magnitude while requiring absolutely no changes to the business intelligence applications. The natural question is "So what has changed? How did they do it?"

The primary difference between mid 80's and now is several fold: database technology has matured; database interfaces (ODBC, JDBC, SQL) are now well defined and stable; powerful components are readily available for open systems; and, very importantly, computer science has come a long way in terms of understanding how to break a query down for efficient parallel processing. Netezza is taking advantage of this convergence and has developed a system not only capable of extracting results from current databases in record time (10X to 100X smaller) but also, for the first time enable databases to grow to terabytes (possibly petabytes) and still deliver query response in "real time" (we define "real time" as the amount of time a query can take and still allow effective decisions to made in a timeframe that is acceptable to a company. For instance a five minute query time may be "real time" for a corporate inventory decision but will probably not be for online retail).

The architecture used by Netezza is a combination of SMP (for the host processor) and MPP (for query engines, called Snippet Processor Units or SPUs). The concept is to break the query into smaller units and let hundreds of SPUs attack the problem. The SPUs are designed to have their own CPU/memory and storage. The proximity of the SPUs to the storage means a lot of filtration can be done at the source and only relevant data (columns for instance) is brought into the memory of the SPUs. The Host processor then combines the results from the SPUs and returns the overall result to the client application. SPUs can be added for scalability of both storage and processing but always kept in balance. The interfaces are ODBC or JDBC and queries use the SQL standard. The BI applications require no changes and can continue to run as before. In essence the Netezza system replaces the general-purpose servers, database servers, database software and NAS or SAN-based storage. The architecture is designed to handle multiple queries from multiple users concurrently. A specially programmed FPGA, is part of the SPU and implements the query function in silicon. Otherwise, the entire system is built from standard off the shelf parts. The host processor runs Linux. The SPUs and the host are connected via gigabit switches, as are the clients.

The system can isolate ongoing transaction updates from affecting the queries. The SPU hardware includes compression/decompression to minimize disk requirements. The catalogs, metadata and transaction logs are RAID protected on the host. Table data is mirrored on the SPUs so a node failure causes no interruption. The architecture is designed to minimize the number of DBAs necessary to run the system. Most work done by DBAs is automated thereby requiring fewer of them.

Product Details

The Netezza Performance System (NPS) is available in three models. The entry model, NPS8100, list priced at $622,000, consists of a single 40U rack with a dual 900 MHz PIII-based host processor, eight Storage Processor Arrays (SPAs), each consisting of 14 SPUs,

for a total of 112 SPUs. Each SPU is complete with a processor, FPGA, real-time OS, and 560GB ATA-based storage. The total raw storage capacity is 4.5TB. The Host processor runs Red Hat Linux 7.1; SQL, ODBC and JDBC are supported and the system is SQL-92 compliant.

The NPS Model 8200 doubles the numbers of racks and the SPUs, for a total capacity of 9.0TB and 224 SPU processors. It is list priced at $1.25M.

The top-end model, NPS 8400, doubles the numbers of racks (over model 8200) to four and SPUs to 448, for a total capacity of 18.0TB. It is list priced at $2.5M.

These models achieve 99.9% uptime via reliability built in at different levels. SPUs are always mirrored. In multiple rack systems, the mirroring is across racks so a complete rack can fail without bringing the system down. SPAs are always added in pairs. Each rack is supported by a UPS, an Ethernet switch and four hot-spare SPUs. The UPS and the Ethernet switch are not single points of failure, except in a single rack system. The only single point of failure is the Host unit itself.

Data integrity is maintained via fully ACID (Atomic, Consistent, Isolated, Durable) transactions; Online Recovery (users can access the system at partial (50%) performance while the recovery happens in the background); Multi-versioning Snapshots (when a transaction is updated, the old version is kept) and Online Incremental Backup and Restore (ADSM compatible).

What Does ESG Think?

We rarely see a system that has the potential for changing how business is done or decisions are made. The timing for this system is impeccable. Data growth is astronomical. The current systems are maxed out. Yes, clusters help. Yes, faster processors help. Yes, InfiniBand will help. But if you are looking for 1) orders of magnitude reduction in time for queries, 2) your databases are (or will soon be) in multi-terabyte range and 3) your business decisions require “real time” factual information, today you will need to go to a purpose-built appliance like Netezza’s.

While it is too early to declare Netezza a winner, we think if they deliver what they promise, they will be alone with such a product for quite a while. One of Netezza’s beta customers was Epsilon, a company that provides marketing services. They experienced a 3 ½-fold reduction in load times and a 50-fold reduction in Update Joins on a 150GB database and have been in awe of NPS’ capabilities. We think businesses in the financial, healthcare, pharmaceutical; retail and several other industries will find this system to be an unexpected gift.

Bottom Line

We rarely make such heavy endorsements. But then Netezza has the potential to alter the way we gather information and make decisions. If they deliver on their promise of performance, scalability, transparence and reliability, they will have a line of customers outside their building.