

DMMRReview

Turning Data Into Intelligence

June 2005

www.dmreview.com



Executive Interview

Challenging the Status Quo

*Jit Saxena
CEO & Cofounder
Netezza Corporation*



NETEZZA

The Power to Question Everything™

www.netezza.com

Challenging the Status Quo

By Jean Schauer

What do you get when you combine a database, storage and analytics? According to Jit Saxena, CEO and Cofounder of Netezza Corporation, you get results from your data warehouse that you thought were impossible to achieve.

A Disruptive Value Proposition

"When Foster Hinshaw and I founded Netezza in September of 2000, the rationale was really all about bringing a disruptive value proposition to what we saw as a major problem in the marketplace," says Saxena. "The problem was that in today's world, you have almost an infinite capability to collect data about all aspects of your business. This can be data about your customers, your suppliers or some of your most fundamental operations issues."


The Netezza cofounders recognized that with the traditional data warehouse architectures, this data could not be cost-effectively stored and analyzed, and they set out to challenge the traditional approach of deploying general-purpose servers, storage and databases for data warehousing.

"We knew that if people could understand their rapidly changing data at a very, very detailed level, mine it and analyze it very frequently or even in a real-time fashion, they could get an edge over any of their other competitors. The architecture we designed allows enterprises to very rapidly and frequently analyze huge amounts of data at a low purchase price and with a very, very low cost of ownership. We were confident that this type of an architecture would be well received, and indeed the reception from the marketplace during the last two-and-one-half years has been tremendous," says Saxena.

The solution Netezza developed to address the marketplace need is the Netezza Performance Server system [NPS]. There are several models that are differentiated based on the capacity that they support. Designed as an appliance that integrates database, server and storage, the NPS system stores, filters and processes terabytes of data within a single unit, analyzing only the relevant information for each query. According to Saxena, the Netezza Performance Server system delivers 10 to 50 times the performance of traditional data warehouse systems at half the cost.

Enterprise-Class Data Warehousing Appliances

"We see ourselves as a supplier of enterprise-class data warehousing appliances. The fundamental architectural concept of our family of products is to put a tremendous amount of intelligence right next to storage," says Saxena. "This is a major departure from the way general-purpose systems have been architected where storage is separate from servers, and the database software is layered on top of that. With that approach, you are basically dealing with three silos. The customer then has the responsibility for bringing it all together and making it work in an optimal fashion. With our architecture, the major performance gains and the major ease-of-use features are the result of integrating all three of these components in a unique architecture that has huge amounts of intelligent storage and allows software engineers to do most of the database operations in a streaming



Jit Saxena
CEO and Cofounder
Netezza Corporation

fashion right next to storage. All the unnecessary data movement from storage to servers to database software is reduced by two orders of magnitude. We call this Intelligent Query Streaming design," he says.

Why is Netezza able to provide increased performance at a lower cost? "The other innovation that we brought to market," explains Saxena, "is that we present this product as an appliance. Rather than requiring our customers to purchase separate storage, servers and software, our appliance has one cable that has a standards conforming interface so that the business intelligence [BI] applications - such as MicroStrategy, BusinessObjects or Cognos - continue to run against the box without any problem. Our customers just plug it in, load the data and they're ready to run their BI applications. Every piece of storage in our system has a dedicated FPGA [field-programmable gate array] on what we call a Snippet Processing Unit [SPU]. Half of each SPU is storage and half is the logic. The SPUs are connected to each other in a massively parallel fashion. For example, our current high-end box, the 8650, has approximately 650 SPUs connected to each other. We've recently announced larger models for up to 100TB of data, which have close to 900 SPUs. That's massive amounts of compute power dedicated to each storage level. When a query comes into the system, the processors work on the query in parallel. They take their portion of the data, analyze it and return the results. The results are then summed up and returned to the application. This reduces unnecessary data movement, and that's why our performance is so much better than anything else out there," he says.

Increasing Data Volumes

As Saxena visits with potential customers, he repeatedly hears the same concern. "The major issue is that as data size increases, query performance and load performance decrease, and cost of ownership increases because the organizations must continuously tune their system to somehow try to maintain their previous levels of performance. The constant upgrading of hardware as well as the constant and ever-increasing investment in DBA resources - often without experiencing any improvement in performance - is a major problem. Companies keep adding DBAs or often require their DBAs to work on weekends. That's a tremendous cost," states Saxena. "Then, when the data size increases again, the whole cycle repeats itself. Until Netezza came on the scene, this was the way most companies dealt with the problem," he adds.

"Whether it is a call data record problem, a clickstream analysis problem, analysis of Web site page search phenomena, analysis of a retail marketing promotion, or analyzing the effect of markdowns, all of these and many other similar data-intensive processes are becoming key drivers of business and require frequent analysis of huge amounts of data. The existing products either cannot do it, or if they do, they do it at a very high initial cost

▶ Jit Saxena

Last book read: *Golf is Not a Game of Perfect*, by Dr. Bob Rotella
Most memorable movie: *Schindler's List*
Favorite entertainer: Jerry Seinfeld
Favorite game: Golf
Dream car: Audi A8 W12
Favorite vacation spot: Loon Mountain, NH
Most admired individual: Gandhi
Favorite TV show: Seinfeld
Favorite sports team: New England Patriots
Hobby: Golf
Favorite type of music: Indian Classical
Favorite food: Hot curries
Favorite beverage: Pinot Noir



▶ Netezza

Year Founded: 2000
Privately Held
Number of Employees: 170
Number of Customers: 28

with a very high cost of ownership. We are addressing that problem head on," emphasizes Saxena.

The Challenger's Challenge

It's not an easy road when you are challenging the status quo. As Saxena explains, "Clearly, we are the smallest and we are the newest. Many people have not even heard of us. But, it keeps getting easier as we get more high-profile, referenceable customers such as Amazon, Ahold USA, Cingular, Orange in the UK, TJX and Shoppers Drug Mart."

"The second obstacle that we encounter," continues Saxena, "is that people just don't believe our claims. In a way, our claims may seem outrageous - I mean, how many times does something come around that is an order of magnitude faster and a fraction of the cost of the existing system? It just doesn't happen very often; however, our value proposition is so strong that companies want to take a closer look. To enable that closer look, we will conduct a proof of concept at our offices in Framingham or at the customer site where we take their data, their applications, their data model and their schema, put it on our box and show them the results. That enables them to validate our claim that they don't need to change any of their applications or any of their data models. They load the data, run their applications, test the interoperability - and then," notes Saxena, "they usually buy it. So far, our track record on that has been close to 100 percent."

A New Paradigm

"We have introduced a new paradigm for data warehousing, and it has huge benefits for the customers," says Saxena. "They are able to do things that they could not do before, at a lower price and a much lower cost of ownership. That is relevant to many, many enterprises as the problem of data continues to get bigger and bigger. Rather than shrug their shoulders and say there is no solution to the data problem, I think more companies should take a look at what Netezza can do." 