



CASE STUDY DATA WAREHOUSE

Data analytics pays dividends

The New York Stock Exchange is deriving more value from its data after replacing a cumbersome warehousing system with an analytics appliance

David Neal

The New York Stock Exchange (NYSE) is the largest financial market in the world in terms of trading volume, listing companies with a total global market value that stood at more than \$27 trillion at the end of April.

Facilitating share dealing on behalf of nearly 3,000 listed companies obviously generates an enormous amount of data, and the NYSE is under huge regulatory and commercial pressure to ensure that this information is managed efficiently and securely.

Incredibly, however, data management was until recently a real problem area for the organisation.

Accessing and analysing information held in its archives was taking so long that the NYSE struggled to get any value from it

It would take us until noon the next day to load our new data onto the old system

at all, according to chief architect, Brian Clarke.

The organisation decided it had to have a new data warehouse, eventually opting for a solution from Netezza.

“Our old platform was a data warehouse, but although our application tier was working



The NYSE has cut the time it takes to load and extract data from hours to seconds

well, we were failing at the data level. In fact it was taking as long as 12 hours to load the data,” Clarke said.

This lag would be unacceptable in almost any industry these days, but in one as fast-paced and tightly regulated as the financial markets it was unforgivable.

And this was not the only usability issue with the system that the NYSE faced.

Not only were requests taking a long time to fulfil, but loading data into the system was also problematic. This led the NYSE to become so disillusioned with the capability of the technology that it began to question whether any data warehouse from any vendor could meet its needs.

“It would take us until noon the next day to load our new data onto the system, and when you are dealing with timeframes like this there is the chance that you will never catch up,” said Clarke. “The old system was also very difficult from an administration point of view. It was a

massive undertaking every time we wanted to perform a task. Any queries that were put through the system were running very slowly, and even working with the vendor that supplied the kit to us, we couldn’t solve these issues. Because of this, we didn’t believe that any other available solution would do any better.”

The NYSE first approached Netezza in 2006. Following these initial talks, Clarke looked at two other systems from competing firms, but decided to work with Netezza because of the maturity of both the company and its flagship Netezza Performance Server solution. Within weeks of signing a deal the system was up and running.

According to Netezza’s director of product marketing, Phil Francisco, Performance Server can deliver 10 to 100 times the performance of traditional data warehouse systems thanks to an architecture that combines a relational database, server and storage with patented streaming



technology. This technology “puts processing power right at the source, next to the data, so it can be processed ‘on stream’,” he added.

The decision to go with Netezza has worked out well for the NYSE. “We only had one shot to get this right,” Clarke said. “I expected it to work well, but I was surprised at how easy it was to integrate the system and get it to work with our data model. The support we got from the vendor was critical to ensuring this smooth integration.”

The difference in performance compared with the old system has been remarkable, Clarke said. “If we had had to load the same amount of data into our old system from a well known software giant, it would have taken many, many months. The new system is so ‘data ready’ that we are finding it possible to use data in more useful ways,” he said.

According to Clarke, data can now be used to perform a variety of functions that were highly problematic with the old system.

These include tracking the value of a listed company, performing trends analysis and searching for evidence of fraudulent activity.

With the old system, any sort of query would take upwards of six hours to fulfil, but with the

Netezza Performance Server, the average time it takes to complete a request is 20 seconds. “So far,” Clarke added, “we have not come across a query that we have not been able to optimise.”

The new data warehouse has also had a huge impact on the NYSE’s

ability to comply with the numerous financial regulations.

“Regulatory requirements mean that we have to keep data for seven years, and we have to be able to access it and provide it on request,” Clarke said. “Now we can meet requests on the same day that they

are issued.” Francisco explained that the data compression technique used in the firm’s technology enables firms to manage their data in a much more efficient way, adding that it can de-compress data as fast as it can compress it.

> www.netezza.com

Q&A: Data boost

■ What was the business need?

The New York Stock Exchange (NYSE) wanted a new data warehousing solution to enable it to get more value from its data and speed up many of its processes.

■ What was the chosen solution?

The company opted for a data warehouse and analytics solution from Netezza, called the Netezza Performance Server.

■ Benefits so far?

The time it takes to back up and access data has decreased dramatically, greatly helping the NYSE to meet its compliance obligations. The system also allows it to put data to more uses, such as fraud detection.