



New York Stock Exchange datawarehouse slashes search times

Off-the-shelf system speeds up NYSE data

Karl Flinders

karl.flinders@rbi.co.uk

The New York Stock Exchange (NYSE) is using off-the-shelf technology to cut the time taken to access business-critical data on its network from up to 26 hours to just over two minutes.

The project will help the NYSE better compete with organisations such as the London Stock Exchange in selling share trading data.

The NYSE, which runs stock exchanges in Europe, the US and Asia, has replaced an Oracle 10 relational database with datawarehousing appliances from Netezza, allowing it to conduct rapid searches of 650Tbytes of data.

The NYSE's project is part of a wider trend for financial services to move away from complex legacy technology towards off-the-shelf products, which give businesses greater flexibility to meet changing market demands.

The NYSE technology, which was installed in December, collects, stores and manages trading data from the New York exchange, its Paris operation Euronext and its

KEY POINTS

- ▶ Time taken to access business data on network cut from hours to minutes
- ▶ Relational databases replaced with Netezza datawarehouse appliances
- ▶ Search engine and storage in the same unit; multiple chips speed search
- ▶ Trend for financial services to move from legacy to off-the-shelf products

Asian exchange Archipelago.

The datawarehousing appliances combine a database search engine and storage in the same unit and use multiple processors in parallel to find the data quickly.

Brian Clark, chief architect at the NYSE, said, "It was painful to put data in and get it out with the previous system."

The previous system had to send large quantities of mostly irrelevant data from storage devices to the database over the network to find the required information.

Clark said the appliances, which carry out server, storage and database management activities, will pay for themselves on the storage benefits alone.

"The financial analysis we did was just on storage volume, and we were able to get a return on investment in one year," he said.

He said the Oracle system used

too much storage capacity because IT staff had to build indexes to help search for information. This meant some data was repeated.

Ralph Silva, an analyst at TowerGroup, said if the exchange can source data faster, it can cut the costs associated with answering questions from the regulators.

"If you have good data management, you can answer the regulator's questions in seconds without the need to put three or four people on the job," he said.

Silva said the NYSE would also be able to sell more data. "The big market for data is from the algorithm writers, who use it to test algorithms against [historical market conditions] to see if they work."

Stock exchanges face IT challenge
→ computerweekly.com/226952.htm

Nationwide off-the-shelf, p10