

HP Oracle Database Machine, New Approach to Data Warehousing

By Drew Robb

June 17, 2009

Data warehousing has become even more important recently as companies seek to store ever-larger and more embracing repositories of information. So it makes sense that instead of simply having data warehousing software that sits on top of myriad disparate systems and tries to make sense of it, that eventually the vendors would get together and create an entire data warehousing ecosystem from the ground up.

Enter the HP Oracle Database Machine, a new approach to data warehousing that leverages 20 Gb/s InfiniBand by Voltaire (Billerica, Mass.). By leveraging InfiniBand as the server-to-storage interconnect instead of Fibre Channel (FC), significant performance gains can be realized.

"The HP Oracle Database Machine is a new approach to data warehousing that leverages 20 Gb/s Voltaire InfiniBand as the server-to-storage interconnect instead of Fibre Channel resulting in extreme performance gains," said Asaf Somekh, vice president of marketing at Voltaire. "It is designed to replace legacy data warehouse systems and help businesses that rely on massive amounts of data to process it faster and more efficiently. The HP Oracle Database Machine can also connect to legacy storage and seamlessly load data and process it faster."

At the heart of this system, is the HP Oracle Exadata Storage Server, which has smart storage software that offloads data-intensive query processing from database servers closer to the data. This results in much less data getting sent over fast InfiniBand interconnects, improving both query performance and concurrency.

The HP Oracle Database Machine is a complete system, including software, servers, networking and storage, designed to run multi-terabyte data warehouses at least 10 times faster than conventional data warehouse systems, according to Voltaire. On the software side, it runs Oracle Database 11g and Real Application Clusters on Oracle Enterprise Linux.

Oracle Exadata is a family of high-performance storage software and hardware. It combines software from Oracle and servers from HP. It takes advantage of massively parallel architecture to increase data bandwidth between the database server and storage. The Oracle software offloads data-intensive query processing from Oracle Database 11g Servers and does the query processing closer to the data.

The HP Oracle Exadata Storage Server is based on HP ProLiant DL180 G5 server building blocks and can provide up to 12 TB of storage capacity. Oracle Exadata software comes pre-installed. In addition to extremely fast query processing for your large data warehouses, the massively parallel architecture offers linear scalability and mission-critical reliability.

Voltaire provides the entire InfiniBand infrastructure. Its 20 Gb/s InfiniBand switches provide high-speed connectivity between the Exadata storage and Oracle database servers. This eliminates network bottlenecks that could limit query performance.

Who are the customers?

Who is using it? Somekh gives the example of LGR Telecommunications (Atlanta), a provider of data warehousing, analysis and reporting solutions for global telecommunications companies. It uses the HP Oracle Database Machine to help its clients process data.

"Today's mobile phone and other communications devices create massive amounts of data on wireless networks," said Somekh. "Managing this data effectively is becoming exponentially complex."

Customer call data records (CDRs), for example, are produced every time individuals place a call. These have increased from 10 CDRs per day to 100 CDRs per day due to new devices such as iPhones. Multiply that by all the wireless users around the world and this translates into billions of CDRs for wireless carriers to manage and store.

"LGR can deploy their system on any telecommunications network around the world, take these huge volumes of data, and transform them into a single consistent format, regardless of the type of network," said Somekh. "By improving storage and management with Voltaire and Exadata, LGR is helping customers gain new insight more quickly for important activities such as business intelligence, fraud detection, and network performance monitoring."

By replacing the FC link between the host database and storage with Voltaire InfiniBand in the HP Oracle Database Machine, LGR has experienced major performance improvements. According to Somekh, queries that used to take minutes now take seconds. He cites benchmark tests comparing the new system to the previous data warehousing system.

"LGR saw performance improvements between 20x and 60x," he said. "Queries that used to run for more than 30 minutes were completed in 30 seconds. Other benefits include faster installation at LGR's customers (i.e., installation in weeks rather than months), which leads to quicker ROI."

<http://www.enterpriseitplanet.com/storage/features/article.php/3825606>