Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.935.4445 F.508.988.7881 www.healthindustry-insights.com

NEWSFLASH

Oracle Unveils Its New Exadata Database Server at Oracle OpenWorld 2008

Healthcare Provider IT Strategies
Healthcare Payer IT Strategies
Life Science Business Systems Strategy

#HI214522

Marc Holland
Lynne A. Dunbrack
Alan S. Louie, Ph.D.

NEWS UPDATE

This Health Industry Insights NewsFlash discusses the announcements made during the recent Oracle OpenWorld event. In his Oracle OpenWorld keynote address on Wednesday, October 24, Oracle Chairman Larry Ellison announced the Exadata Storage Server, Oracle's new high-performance storage server and the heart of the new HP Oracle Database Machine, designed specifically for storing and querying huge data warehouses.

Designed by Oracle and built to its specifications by HP, the Database Machine promises to dramatically improved price/performance over past configurations proposed by Oracle to operate similarly sized data warehouses, as well as database machines designed for similar purposes from competitors such as Teradata and Netezza.

OUR VIEW

The Bottom Line

The keynote addresses at Oracle OpenWorld are well known as the stage for blockbuster announcements and this year's were no exception. Ellison opened his presentation by citing forecast data from the WinterCorp, Waltham, Massachusetts, which estimates that the current average-sized data warehouse is now 300TB, and expected to triple in size every two years.

Cost and performance continue to be two of the primary barriers to successful data warehouse deployments. According to Oracle, on most general purpose database servers, performance slows dramatically when the size of the underlying source database exceeds 1TB. The principal reasons are data transfer rates and the preprocessing of the raw source data to prepare queries. Oracle's solution is to offer a machine with bigger "pipes," offer more of them, and offload the preprocessing steps into the primary storage server. These include the scanning, filtering, and matching of row/column selection for many SQL "where" clauses from the database server to the storage server. As a result, far less data will travel from the storage server to the database server, and travel over more and bigger pipes.

Many larger organizations, with experience in the deployment of massive data warehouses, have turned to proprietary database machines, including those from Teradata and Netezza. It is these vendors that are now in Oracle's crosshairs.

The base Exadata server comes equipped with dual quad-core Intel Xeon 2.66GHz processors, 8GB memory, 12 high-performance 300GB or 1TB disk drives, Oracle Linux, and Oracle's parallel query database software. Up to a maximum of 168 high-performance drives ranging in size from 300GB SAS to 1TB SATA drives can be supported in a single complex. Each of the database servers also comes with dual quad-core Xeon processors for a total of 64 CPUs; each of the 14 Exadata servers comes equipped with dual quad-core processors for a total of 112 CPUs providing 14GBps total bandwidth. Servers are rackmounted and each server is equipped with dual InfiniBand busses, with each pair capable of a 1GBps data transfer rate.

With Oracle's parallel query software and an Exadata server, source data can be preprocessed and presented to the data warehouse as prebuilt query tables. While initial shipments will only support database servers running Linux, support for other database operating systems is promised in the coming months.

The three case studies from early adopters that were presented by Ellison, which included a European telecomm provider, a United States–based supermarket chain, and a United States–based commodities exchange, cited performance improvements ranging from 10 to 72x. Ellison also presented a cost comparison that claimed a fully configured Database Machine complex could be delivered for one-third to one-half the cost per terabyte of storage versus comparable machines from Teradata and Netezza. A significant portion of the performance improvement could be attributed to Database Machine's preprocessing capabilities.

Essential Guidance

Healthcare is one of the fastest growing vertical markets as measured by the compound annual growth rate (CAGR) of its IT spending; this trend is expected to continue well into the next decade. Although each of the three industry segments, as defined by IDC, will grow at differing rates, all have and will continue to grow faster than the average overall IT spending across all industries. These segments include healthcare providers, healthcare payers (i.e., health plans), and pharmaceutical and life science firms. Spending on health-related IT projects by government agencies, particularly at the state and local levels, is also expected to rank among the highest of any vertical market.

Data warehousing, and the business intelligence (BI) applications these huge data stores are designed to support, make up one of the fastest growing segments of IT across a wide range of vertical markets, including healthcare. While provider organizations are just beginning this journey, large payers, big pharma, and the federal regulatory agencies that oversee and finance U.S. healthcare are well down the road. Each of these players is investing in IT at unprecedented rates and with the resulting increase in their underlying application portfolios, the increased volume of available transaction data is creating a demand to build richer, and larger, data warehouses to mine the value inherent in this data. For example, in 2006 CS announced a major contract with the Blue Cross Blue Shield (BCBS) Association to build the Blue Health Intelligence data warehouse. The warehouse was planned to accommodate claims data for 79 million members of 20 BCBS plans and, at the time of the announcement, was touted as the largest data warehouse in the U.S. healthcare market.

The hospital IT market, despite the saturation we are seeing at the level of core transaction processing applications, still has a lot of room to grow. As hospitals gain experience with their electronic medical records (EMRs) and fully populate their application portfolios at the transaction processing level, they are preparing to usher in the next generation of IT. We expect that this generation will include a heavy emphasis on business and clinical intelligence applications. These applications will be focused on improving clinical, financial, and operational performance; treatment outcomes; customer service experience; and clinical decision making.

Beyond the hospital market, integration of care delivery processes and data with clinical research, drug discovery, and clinical trials processes presents another opportunity to break down the traditional walls that exist between the major players in the healthcare ecosystem. We expect that these enhanced IT capabilities will drive breakthrough improvements in the quality of clinical quality care and efficiency of healthcare service delivery.

Finally, as more experience is gained with the various models of health information exchange, more widespread adoption of the successful models is expected. One result of this will be consolidation of data from smaller, less successful exchanges.

The net effect of these trends will be to dramatically increase the industry's demand for not only storage and tools to manage these vast amounts of data but also for better technologies to mine it. BI will be a critical tool in the IT arsenal of healthcare organizations. Health Industry Insights is projecting CAGRs for BI tools and related technologies to be in excess of 10% by providers and 12% for payers. To maximize their return on their investment in BI, it will be important for the healthcare industry to draw upon the experience gained by other vertical markets that have more experience with large-scale data warehouses.

If the cost and performance advantages touted by Oracle are verified in future deployments, Oracle's new Exadata machines are likely to be well represented in future healthcare data warehouse deployments. Exciting times, indeed.

LEARN MORE

Related Research

- *Healthcare Provider Business Intelligence 2008: What Can We Learn from the Early Adopters?* (forthcoming)
- *The Practical Faces of Translational Research: Dana-Farber Leads the Way* (Health Industry Insights #HI213727, August 2008)
- *Healthcare Payer Business Intelligence Solution Evolution, 2008–2010* (Health Industry Insights #HI210376, February 2008)
- *Healthcare Payer BI Maturity Model, 2008–2010* (Health Industry Insights #HI210218, January 2008)
- *U.S. Healthcare Payer 2008 Top 10 Predictions* (Health Industry Insights #HI210367, January 2008)
- *U.S. Health Industry Provider 2008 Top 10 Predictions: The Use of IT in the Provider Landscape* (Health Industry Insights #HI210302, January 2008)
- *IBM Showcases Advanced Data Storage and Management Technologies at Its Healthcare Data Summit 2007* (Health Industry Insights #HI208927, September 2007)
- *High-Performance Translational Medicine: Streamlining the Researcher Experience* (Health Industry Insights #HI208454, August 2007)
- *A Vision for Sustainable Intelligence in Integrated Clinical IT Systems* (Health Industry Insights #HI206165, March 2007)

Copyright Notice