



LGR Test Drives HP Oracle Database Machine

Charles Babcock
Sep 25, 2008

Oracle CEO Larry Ellison announced a high-end, high-performance HP Oracle Database Machine in San Francisco yesterday. LGR Telecommunications has been working with such a Database Machine for more than a year, while the kinks were still being ironed out. Chief architect Hannes van Rooyen says it's an ideal design for taking on massive data loads with rapid cycle times.

Van Rooyen concedes LGR, a supplier of custom telecom data warehouses, hasn't pushed an HP Oracle Database Machine out to one of its telecommunications customers yet. But its simulated production loads on the machine and found it handles them. To LGR, a production data warehouse should be built on the scale of the one it supplied to the largest U.S. telecommunications company -- you should be able to figure out who that is, but LGR isn't allowed to name it. It uses two 310-TB data warehouses, and needs to cycle lots of data back and forth between them. That's what the Database Machine is designed to do.

LGR won't be spending more time on beta tests. Now that the machines have become available, it plans to install at customer sites in the next few months, including that large customer in the U.S.

Deploying a Database Machine, said van Rooyen, 28, is a lot faster than deploying additional servers and SAN arrays as your data warehouse grows. LGR ought to know. It's built a lot of standard data warehouses on high-end hardware, such as the HP Superdome, throwing disk arrays at the arrangement. "You have to run the host bus adapters (on the database servers) to the SAN array, reconfigure and rebalance the load," he says. It's a lot of skilled labor.

In one case, "we cycle 260 terabytes of data through the data warehouse every 90 days. There's a lot of velocity to our data," van Rooyen said in an interview at the Moscone Center as Oracle OpenWorld wound down.

Telecom companies are trying to capture who's using what services where, what combinations of services, how frequently, and where the profit lies

in combining services. To do so they need to deal with billions of small records after capturing and loading them into the data warehouse.

Van Rooyen said LGR is planning systems that can handle up to two petabytes of data. Parallel processing is a key to query performance against such large amounts of data, and the Oracle system does a good job of breaking queries down into parallel parts. "We push the core functionality of Oracle" to make use of its query optimization and use of Real Application Clusters, or a clustered version of one database on multiple servers.

Beyond that, LGR brings its expertise to bear to exploit the capacities of multicore chips. Big databases and complex queries are a problem for which multicore processors are the answer. Because both SQL queries and databases can be broken down into discrete parts, they lend themselves to parallel processing techniques better than many business applications, with their serial logic.

IBM, Netezza, Teradata, and others all have their uses of parallelism in their high-end data warehouse designs. But as van Rooyen points out, using it on standard hardware -- the database machine, in effect, is a small grid of HP Proliant servers -- means LGR can produce telecom data warehouses, with their voracious appetite for CPU resources, at a much lower cost than in the past. The Oracle machine retails for \$4,000 a terabyte, he noted.

Van Rooyen says HP may sell fewer Superdomes now that the database machine is available, but it may sell more database machines. "Now when we need more storage, we plug in more Exadata Storage Servers. If we need CPU, who knows, Intel is going to six-core processors. Maybe we'll slide one blade out and another blade in," he said.

In one of LGR's installations, the data warehouse routinely consumes 99% of the CPU power available, and six-core CPUs will be welcome when they reach the site on new HP blades.