An Inside Look at Oracle's (and HP's) New Database in a Box and the 'Accelerator'

– Eric Lai, September 29, 2008

Oracle Corp.'s annual OpenWorld show is usually a showcase for its enterprise software. This year, however, it was all about hardware, as CEO Larry Ellison introduced a new family of database/storage products on Wednesday that it had been working on with partner Hewlett-Packard Co. for three years. Here's an FAQ about HP-Oracle's Database Machine and Exadata Storage Server:

Which one is the "database accelerator" that Ellison had been teasing us about?

The Exadata Storage Server, a standard rack-mountable HP ProLiant DL180 G5 server sporting two Intel quad-core CPUs connected to 12 hard drives of 1TB each.

What makes it different than a typical Linux-based storage server is the fast parallel-query software built into it, which allows the Exadata to perform a number of database functions locally. The Exadata can "optimize queries" by doing lower-level calculations closer to the raw data. Only the results are then sent to the actual database, which aggregates the results in order to perform the final number crunching.

"The intelligence allows us to reduce the amount of data flowing across the interconnect," said Ellison, which, as he rightfully pointed out, is the source of most of the delay in modern database systems.

Besides reducing the data being transmitted through the network, the Servers are also equipped with two InfiniBand connections for high-speed data transfer. Ellison says Exadata users can expect a real-world bandwidth today of 1Gbit/sec, which he claimed is far faster than conventional disk storage arrays.

And the Database Machine?

The imposingly named Database Machine is essentially an all-in-one database and storage solution for companies seeking building blocks as their databases near the petabyte range.

The Database Machine comes on a single rack and includes 14 Exadata Storage Servers, eight Linux servers running Oracle Database 11g Enterprise Edition, a total of 368GB of RAM and 168TB of disk space. Ellison joked, "This holds a lot of songs. It's 1,400 times larger than Apple's largest iPod."

Not only is it massive; Oracle also claims it's ultrafast because the machine brings the database servers even closer to the Exadata servers in a network with a raw I/O of 14Gbit/sec.

How fast is the new breed?

Ellison said the Database Machine was between 10 and 72 times faster than conventional Oracle databases. A number of Oracle customers have been testing the Machine for a year, putting their actual production workloads onto half-sized Oracle Database Machines ("because we're really cheap," quipped Ellison).

Simeon Dimitriov, an IT manager at M-Tel, an Eastern European telephone company, said in a video interview shown by Oracle that queries and reports generated were 28 times faster. It ran Oracle Database on 2 IBM Power 570 servers connected to a midrange EMC CX3-40 storage array.

LGR Telecommunications Inc. CEO Grant Salmon oversees 310TB of caller records on an HP Superdome Server connected to a Hitachi XP24000 storage array. He said queries involving caller data records that used to take 30 minutes took just a minute when run by a half-size Database Machine.

The Chicago Mercantile Exchange's workloads ran between 10 and 15 times faster, and grocer Giant Eagle Inc.'s 5TB sales data warehouse ran 16 times faster, according to Ellison.

What did Oracle have to say about its competitors?

Ellison took pot shots at NCR Corp. spinoff and long-time data warehousing vendor Teradata Corp., as well as Netezza Corp., perhaps the leading member of the new crop of appliance vendors. Ellison claimed that the Database Machine is faster than a five-rack Teradata 5550 cluster, which he said lacks Oracle's new smart-query features and thus is bogged down by the need to transfer much more data from the storage array to the databases. "Architecturally, [the Teradata 5550] just can't compete," he said.

Netezza has a newer architecture, Ellison said, and thus, like Oracle, it has both big bandwidth and intelligence in its storage servers. The difference? "Ours runs Oracle; theirs does not," he said. Netezza's "overall database capability is very primitive," he said, citing a feature, B-Tree Indexing, that Netezza supposedly lacks. Ellison joked, "Even I studied about it in school, and you know how long ago that was."

Teradata provided a measured reply. But Netezza's president blasted the Oracle-HP products for not being designed "from the ground up" by engineers in the same company, and instead being patched together "with glue and spit."

Ellison, uncharacteristically, didn't attack Microsoft Corp., which leapt into the very-large data warehousing market earlier this year when it bought appliance vendor Datallegro Inc.

How much will Oracle's products set you back?

Oracle wasn't clear about how much a single Exadata would cost. But a fully-loaded Oracle Database Machine lists for US$2.33 million, including the hardware and all of the databases licenses.

If you can use 168TB of storage, then Oracle may actually be the low-cost option (for now), working out to $14,000 per terabyte. According to Ellison, Teradata's closest-equivalent system costs $35,000 per terabyte, while Netezza charges about $29,000 per terabyte for its 10100 system. Even Greenplum charges $20,000 per terabyte for its systems, according to Yara.

What do experts and users think about the new gear?

"Oracle Exadata Release 1 is hardly going to put Teradata, Netezza or Greenplum out of business," wrote analyst Curt Monash. "Medium [to] long term, the Exadata technical strategy could work very well. Exadata storage management addresses some of the problems with shared-everything. Oracle RAC (Real Application Clusters) addresses other, and it may not take many releases before Oracle gets query parallelization right as well."

Database consultant Peter Scott enthused in his blog, writing, "For strange people like me, people that see the world as moving large amounts of data around, it was exciting news."

Christo Kutrovsky, a database administrator at The Pythian Group Inc., thinks the biggest plus of the Database Machine is the fact that it comes preconfigured and pretested for top performance. "Ninety percent of the problems I've seen are due to improperly configured systems," he wrote. "Using Exadata necessarily and immediately solves all of these issues."

"From the DBA perspective, it's heaven," Kutrovsky wrote. "No arguing with storage people for dedicated spindles, no arguing with CIOs about big vs. small disks, no arguing with systems administrators for ASM [automatic storage management]. No hiring of expensive consultants to 'tune' the system or apply best practices. You may laugh at all of the above issues, but many shops are exactly like that. Especially the big ones (the target market for Exadata), where everyone is too afraid to change anything in case they get blamed if it doesn't work. The 'best practices' are the only practices with the Database Machine."

Will the Database Machine or Exadata Storage Server ever run operating systems other than Linux?

Yes, said Ellison, though he gave no timetable.