



Lateral Thinking in the VLDB Market

Infrastructure Strategies, Server Infrastructure Strategies

Charlie Garry

Traditionally, very large databases (VLDBs) were scaled out for massively parallel processing (MPP) across clusters of servers and associated disks. The high price point of such systems would stagger all but the largest organizations. Netezza, a newcomer to the VLDB market, currently offers an appliance that promises to make VLDB processing affordable even to the midmarket.

After more than 40 years of commercial data processing, many organizations are shifting emphasis away from transaction processing and devoting more of their efforts toward data analysis. Analyzing large volumes of data efficiently has been a problem that has fascinated computer scientists for years. The answer has always been to take the huge effort of processing terabytes of data and break the task down into a series of smaller parallel jobs. This method was dubbed "massively parallel processing." Vendors such as Teradata, Informix, and IBM have delivered commercial databases engineered specifically for MPP and continue to support many of the world's largest databases.

Although MPP systems have yet to hit a scalability ceiling, their price points (hardware and software) have certainly placed them out of reach for many organizations. For example, our research found a six-node (i.e., server) cluster supporting 4.4TB (2.2TB usable) and database software with a discounted total price of \$2,409,014. Through 2010, we expect the market for data analysis to experience significant growth. This includes traditional data warehousing as well as time-sensitive critical analysis (e.g., usage fees, fraud detection). We believe the data analysis market's VLDB requirements are not limited to Global 2000 (G2000) organizations but are increasingly a fact of life for small and midsized businesses (SMBs). These emerging vendors will offer outsourced data analysis services (either hosted or packaged) to vertical markets. Although the organizations themselves will be small, their data processing requirements will be huge. Turnaround times will be tight, and margins will be crucial. Unfortunately, the leading providers for warehousing and analytic platforms (e.g., Teradata, Oracle, IBM, SAS) often require large capital investments and skilled administrators (e.g., architects, DBAs), making them unsuitable for much of this growing demand in the SMB space.

Enter Netezza, a recent entrant into the VLDB market. Its Netezza Performance Server (NPS) is a rack-mounted data analysis appliance built from the ground up, using a combination of commodity hardware (Intel-based) and open source software (e.g., Red Hat Linux, modified PostgreSQL). What makes Netezza's solution particularly interesting is the innovative way it has engineered the appliance for low cost and high performance. In fact, the least technical aspect of the NPS is its database software. Unlike Oracle — with its myriad of indexing, partitioning, and materialized views — NPS offers near-zero physical-design options. What sets NPS apart is its unique hardware configuration that builds on the previous work in MPP into something Netezza refers to as asymmetrical massively parallel processing. The Netezza appliance has several models (e.g., 8100, 8200, 8400), yet each contains a single two-processor master node. What differentiates each model is the number of snippet processing units (SPUs), which contain a combination of a CPU and a disk drive. Data is hashed across all the SPUs. The CPU on the SPU processes the majority of the requested data. This design enables Netezza to utilize 114 CPUs on its smallest (8100) model and as many as 674 processors on its largest (8650) model. Resulting data from the SPU is forwarded to the master node for final projection.

The Netezza database software is extremely basic. The only access path that will ever be chosen is a full table scan, for which the NPS architecture is uniquely suited. In fact, the database is so simple that indexes are neither required nor used, eliminating the

META Trend: By 2005, XML-based content storage and rendering (across structured and unstructured content types) will become a key DBMS differentiator, along with data movement and transformation. Centralized data warehouse architectures will emphasize database workload prioritization and data archiving to support the increasing need for near-real-time analytics through 2006.

need for sophisticated physical design and SQL tuning. The NPS supports most open connectivity standards like ODBC and JDBC. The database engine is ANSI SQL92-compliant, which means it currently lacks some of the more robust analytic SQL extensions found in other databases. With simplicity come positives and negatives. The positive side is extremely low maintenance (less than one DBA and one SA required) with very fast time to value. On the negative side, the configuration may be suitable only for tens of users versus hundreds due to concurrency issues depending on the workload mix. High numbers of singleton select (specific row) may be an issue, but currently that is not the market Netezza is looking to tackle.

We have been following Netezza for more than a year (pre-general availability) and have had a unique opportunity to speak with virtually every customer and several prospects. The results of that research indicate that Netezza presents a genuine solution. Currently, it has begun to carve out a niche in the telecommunications (telco) market. Here, organizations process billions of call detail records (CDRs) per month. The massive amount of CDR data has made processing all the data virtually impossible. One telco organization we spoke with wanted to accurately capture usage data across its network so that it could correctly bill other telcos for its network use. Because the resulting data processing needs were so large, the organization had to process only odd-numbered months, resulting in millions in lost billings. Daily CDR processing required almost two days to process utilizing a Sybase/Sun E10K platform. After installing the NPS and a simple port of the database schema, daily processing was reduced from 36 hours to one hour. This meant that the telco could now process all its data. With the extra time, users are now able to query the CDR data for even more information than before. However, this was not a typical data warehousing environment, as the detail data is kept only for a few days and summary data for a few months, due to the huge volumes of data and the lack of clear business requirements. Currently, Netezza's storage capacity scales to 27TB of user data, which is large enough for the vast majority of organizations.

We wondered about the reliability of Netezza's appliance and spoke to several customers and prospects that experienced system crashes. In these instances, users reported acceptable turnaround times by Netezza's technical support and field support staff. Usually measured in hours, more benign issues were often resolved within days. The appliance does have some fault tolerance built in. Each SPU is paired with another SPU in the rack for mutual failover. For example, SPU1 contains its own data and is the mirror of SPU4. If SPU4 were to crash, SPU1 would be used to create a new SPU4 from one of the spare SPUs included in the rack. Disk efficiency is comparable to vendors such as Teradata (i.e., 1/3 user data, 1/3 mirror data, 1/3 spool or temp space). We spoke with users that did reliability testing during proof of concepts and found failover capabilities were satisfactory. However, Netezza does lack expertise in areas such as business continuity (e.g., multisite capabilities), hot backup, and disaster recover features (e.g., block-based replication). All this functionality will be delivered as Netezza gains more practical experience in the market.

Price/performance was the most-mentioned reason why early adopters have taken a chance on such a young organization. One user claimed NPS provided a solution that was one-third the cost of the previous Oracle/Sun combination with performance that was orders of magnitude better. In one test, a user executed 400 concurrent queries, each against its Oracle platform and the Netezza platform. The user reported that Oracle finished in 38.5 hours with 22 errors, while the NPS took just 1.5 hours with no errors. Load times were also impressive, as users reported 50GB-60GB per hour without any tweaking of the load utility. Netezza has also posted impressive wins in head-to-head user benchmarks against both Teradata and IBM DB2 UDB for Unix/Windows.

Netezza must demonstrate penetration in other markets such as financial services and retail to gain the confidence of G2000 organizations. In addition, to penetrate the high-end enterprise data warehouse market, the NPS must gain credibility beyond its current niche status. Adding SQL99 analytic extensions is a must — along with improved concurrency — but its unique appliance simplicity should enable Netezza to force its larger competitors into tougher on-site proof of concepts and compete quite well.

Bottom Line

Netezza is taking an innovative approach to tackling perhaps the most important new database market. Its ability to execute and expand into new markets will determine if it can break into a leadership position in the EDW market.

Business Impact: Data analysis capabilities will be the defining feature of successful organizations.