

Vertica positions itself as an analytics platform following growth spurt in 2010

Analyst: Matt Aslett

Among the new breed of analytic database specialists, **Vertica Systems** was always one of the most likely to mount a serious challenge to the established data-warehouse providers. The company now claims to have made that leap following a breakthrough year in 2010, when it grew both revenue and customers by more than 100%. It is now positioning itself as an analytics platform provider with the addition of more user-defined analytic functions.

The 451 Take

It has been some time since we had an update from Vertica, and we are impressed by its growth in the last year. Crossing the 100-customer barrier in September 2009 was a significant signal of its progress, so to have more than trebled that less than eighteen months later shows how the company is continuing to grow rapidly. The technology is also maturing, as is the marketing, and it is interesting to note that what previously differentiated Vertica (columnar, massively parallel platform) is now pretty much par for the course. The emphasis on the Vertica Analytics Platform is reflective of that, as well as wider trends encouraging in-database analytics, although Vertica is by no means the only company looking at in-database analytics as a differentiator.

Context

When Vertica Systems introduced version 4.0 of its Vertica Database in early 2010, it announced its intention to directly challenge the incumbent data-warehouse providers for enterprise data-warehouse (EDW) projects. Founded in 2005, the company had already attracted over 100 customers at that point by positioning itself as a platform for data marts, but was looking to mount a greater challenge for EDW deployments with support for mixed workloads. With the appointment of former **F5 Networks** and **Acopia Networks** executive Chris Lynch as CEO in April 2010, Vertica also stated its intention to grow internationally, with investment in building a direct sales and marketing presence in the major European cities and in training local channel partners and support providers. Although the company does not disclose revenue details, Vertica recently claimed to have grown both its revenue and customers by over 100% in 2010, and it now claims 328 customers, compared to 120 in February 2010. Having previously had just four salespeople in its Massachusetts headquarters, the company now claims 14 sales teams around the world, and has grown its total headcount to 100, up from 80 in February 2010. The company also claims to be profitable.

Products

Vertica's core product is the Vertica Database – a massively parallel processing (MPP) column-oriented database based on the C-Store column-store database project led by database pioneer Mike Stonebraker at **MIT**. The product is now on version 4.1, which was released in November 2010 and focused on an updated driver architecture, as well as cluster-aware resource management, improved SQL support, backup and recovery, and query performance. Version 4.1 also includes improvements to large-scale update and delete performance, which had initially been delivered in version 4.0 and was one of several features in that release – alongside workload and resource management, time series and sessionization support, and improved memory overflow – that had encouraged the company to claim that Vertica was now ready to challenge the EDW incumbents.

Strategy

Vertica Database was already in use in EDW deployments, but the company's strategy had initially been to focus on small data mart deployments to prove its value to a particular customer before being adopted for increasingly large and mission-critical environments. A classic customer adoption example initially involved a 1TB license that grew over time to 20TB and is now 1PB. While the Vertica Database remains at the core of the company's value proposition, it is now positioning it as the Vertica Analytics Platform, with the Vertica Database at the heart of a platform that also includes operational and disaster-recovery capabilities, integration tools and connectors, administration tools, database design tools, and real-time in-database analytics. The latter will be a particular focus in 2011 as the company looks to exploit improved systems performance by bringing analytic processing closer to the data by adding a UDx Framework for user-defined analytic functions (such as scalar, table and aggregates). The UDx Framework is being built from scratch to make use of the Vertica Database's MPP architecture and compression capabilities. The company also plans to deliver the most popular native functions, covering the likes of real-time statistical functions, pattern matching/discovery, graphing and MapReduce type functions, as SQL extensions.

Technology

Since MPP architectures and column storage have become accepted technologies for analytic databases, Vertica expects to spend less time evangelizing the four Cs at the heart of its database technology: column storage, clustering, compression and continuous performance. It will continue to enhance them, however, as it builds out the capabilities in the Analytics Platform. Also on the roadmap are enhancements to the FlexStore column-storage architecture, with improved lifecycle management and a new cache layer, as well as a new Web-based interface and management console. Vertica was one of the first database vendors to realize the complementary value of Apache Hadoop as a data-processing framework for unstructured data, and is building on its Vertica Connector for Hadoop and its partnership with **Cloudera** with a prototype project called SQueal, which is an engine for running Pig programs on both Vertica and Hadoop. Additionally, while Vertica has been at the forefront of enabling analytic databases on virtual environments, it is pushing further forward with Project Cumulus to deliver database bursting across physical, virtual and cloud platforms, as well as federated management and re-provisioning of Vertica Databases wherever they

reside. These new capabilities will begin to be delivered with Vertica version 5.0, which is due in the second quarter.

Competition

When we write of Vertica challenging the incumbent data-warehousing providers, we are referring to **Oracle, IBM, Teradata** and **Microsoft**, which dominate the market for enterprise data warehouses. Other industry heavyweights in the market include **EMC** and **SAP**, thanks to their acquisitions of **Greenplum** and **Sybase**, respectively. Although they lack the market share of the big four, EMC and SAP have the resources to invest to grow their respective database businesses. The next group of vendors, which includes Vertica along with **Aster Data Systems, ParAccel, Infobright, SenSage, Calpont, Kognitio, Ingres, SAND Technology** and **1010data**, among others, lack the comparative resources, but have the benefit of data-analytics expertise and innovative approaches to differentiate themselves. Vertica reports that it most often comes up against **Netezza**, the appliance specialist recently acquired by IBM, while the focus on user-defined functions for in-database analytics would, in our opinion, often see the Vertica Analytics Platform compared to Aster Data, ParAccel and Calpont.

SWOT analysis

Strengths	Weaknesses
Following its rapid growth over the last eighteen months – and consolidation elsewhere – Vertica is now arguably the biggest independent challenger to the data-warehousing giants.	The company's former differentiators – an MPP architecture and columnar storage – are now almost par for the course when it comes to analytic databases.
Opportunities	Threats
The growing volume of data is encouraging enterprises to look to new and emerging data-storage and -processing technologies.	Oracle, IBM, Microsoft and Teradata dominate the market, while EMC and SAP are new entrants. Vertica has the expertise, but does it have the resources to compete long term?

Reproduced by permission of The 451 Group; copyright 2009-10. This report was originally published within The 451 Group's Market Insight Service.

For additional information on The 451 Group or to apply for trial access, go to: www.the451group.com