Magic Quadrant for Data Warehouse Database Management Systems

The data warehouse DBMS market is undergoing a transformation, including many acquisitions, as vendors adapt data warehouses to support the modern business intelligence and analytic workload requirements of users. This document compares 16 vendors to help you find the right one for your needs.

WHAT YOU NEED TO KNOW

Despite a troubled economic environment, the data warehouse database management system (DBMS) market returned to growth in 2010, with smaller vendors gaining in acceptance. As predicted in the previous iteration of this Magic Quadrant, 2010 brought major acquisitions, and several of the smaller vendors, such as Aster Data, Ingres and Vertica, took major strides by addressing specific market needs.

The year also brought major market growth from data warehouse appliance offerings (see Note 1), with both EMC/Greenplum and Microsoft formally introducing appliances, and IBM, Oracle and Teradata broadening their appliance lines with new offerings. Although we believe that much of the growth was due to replacements of aging or performance-constrained data warehouse environments, we also think that the business value of using data warehouses for new applications such as performance management and advanced analytics has driven — and is driving — growth.

All the vendors have stepped up their marketing efforts as the competition has grown. End-user organizations should ignore marketing claims about the applicability and performance capabilities of solutions. Instead, they should base their decisions on customer references and proofs of concept (POCs) to ensure that vendors' claims will hold up in their environments.

Many trends — such as poor data warehouse performance, heavy competition and widely disparate marketing claims — will continue through 2011 and beyond. They will be joined by many new pressures arising from the desire to deliver greater business value with new applications, such as demands for further reductions in latency, the acquisition of appropriate data and greater performance. We describe the new and continuing trends in depth in "The State of Data Warehousing, 2011" and "Data Warehousing Trends for the CIO, 2011-2012."

This Magic Quadrant deals with one of the primary building blocks of data warehouse infrastructure. As such, it should interest anyone involved in defining, purchasing, building and/or managing a data warehouse environment — notably CIOs, chief technology officers, members of business intelligence (BI) competency centers, infrastructure, database and data warehouse architects, database administrators (DBAs) and IT purchasing departments.

Gartner

Note 1. Definition of a Data Warehouse Appliance

A prepackaged or preconfigured, balanced set of hardware (servers, memory, storage and I/O channels), software (operating system, DBMS and management software), service and support, sold as a unit with built-in redundancy for high availability and positioned as a platform for data warehousing. Further, it must be sold on the basis of the amount of SSED ("raw data") to be stored in the data warehouse and not of configuration (for example, the number of servers or storage spindles). Our performance criteria have some flexibility to accommodate vendors that have several variations, based on desired performance SLAs, and the type of workload intended for the appliance. Our primary concern is that the user (buyer) cannot change the configuration due to budgetary issues, thereby adversely affecting the performance of the appliance.

MAGIC QUADRANT

Market Overview

In our previous "Magic Quadrant for Data Warehouse Database Management Systems" we described the market in 2009 and the changes we expected for 2010. We also identified market forces, end-user expectations and vendors' resulting solution approaches as including:

- Increased demand for optimization techniques and performance enhancement.

- The argument made by purchasing departments that buying power increases when dealing with a single, incumbent vendor.

- Prepackaged, prebalanced warehouse environments delivered using data warehouse appliances.

- Expectations for the delivery of on-site POCs.

- Cost controls and data warehouse performance management.

- Demands for delivering a fully mixed workload.

- Demands for departmental analytics delivered quickly via data marts.

- Wider indexing and fast performance within clusters of data, delivered via column-based solutions.

- A wave of new data warehouse implementers seeking fast-track, low-risk delivery.

- Global organizations seeking distributed solutions as potential architecture.

One other factor worth considering is the effect of the economic crisis on the DBMS market in general and specifically on DBMS vendors in 2009 and 2010. In 2009, revenue in the relational DBMS market was flat, at $18.8 billion. Given that most other IT sectors — servers, software, telecommunications and services — were declining, the flat performance of the DBMS market is actually an indication that organizations were appreciating the greater business value derived from applications run on data warehouses. Operational analytics, performance management, operational BI and predictive analytics demonstrate to most organizations enough business value to justify continued investment, even when the economy is depressed. We also believe that revenue returned to growth in 2010 and that it will grow further in 2011. In short, the DBMS market has weathered the economic storm of 2008 and 2009.

Many of the trends mentioned above had an impact on the market in 2010, and some vendors met the resulting demands better than others. In 2010, few new vendors entered the market, and existing vendors released new functionality and platforms and increased their marketing. The most significant market shift came with three acquisitions: Sybase was acquired by SAP, Greenplum by EMC, and Netezza by IBM, the first two acquisitions adding two more large vendors to the market. This shift is not only prompting many questions from customers of the acquired vendors, but also raising the question of whether further consolidation will occur. Also in 2010, Microsoft released its much anticipated SQL Server 2008 R2 Parallel Data Warehouse (PDW), a massively parallel processing (MPP) data warehouse appliance based on SQL Server, and EMC/Greenplum released its first data warehouse appliance.

As in 2009, Gartner clients still report performance-constrained data warehouses during inquiries. Judging from these discussions, we estimate that nearly 70% of data warehouses experience performance constraint issues of various types. These typically affect data warehouses with varying levels of mixed workload (see Note 2), especially those with high query counts, mixed query types, and growing integration with both operational and BI applications. Importantly, performance-constrained warehouses are difficult to identify because the enterprises affected often have not established clear service-level expectations, making it impossible to determine how the warehouse is performing relative to a service-level agreement (SLA).

In 2011 we will be watching many new, small vendors that did not meet the inclusion criteria for this year's Magic Quadrant — for example, Algebraix Data, EnterpriseDB and Exasol. Gartner believes there will be additional consolidation in 2011 as the

megavendors continue to acquire innovative products to increase the functionality and capabilities of their DBMSs. As shown in Figure 1, there are several small, innovative vendors that could reduce development efforts by, and accelerate new functionality for, larger vendors, if acquired by them. Importantly, these small, innovative vendors have a customer base approximately the same size as those of Greenplum and Netezza at the time of their acquisition in 2010.

In addition, we believe the data warehouse DBMS market will continue to change in 2011 in order to fulfill the demand for high speed, lower latency and large volumes of data brought about by new high-value applications. The primary forces that we believe will have an impact on the data warehouse DBMS market in 2011 are:

- Increased demand for optimization techniques and performance enhancement.

- The need for data warehouse infrastructure to manage "extreme data" (see Note 3).

- Increased demand for data warehouse appliances.

Figure 1. Magic Quadrant for Data Warehouse Database Management Systems

Source: Gartner (January 2011)   As of January 2011
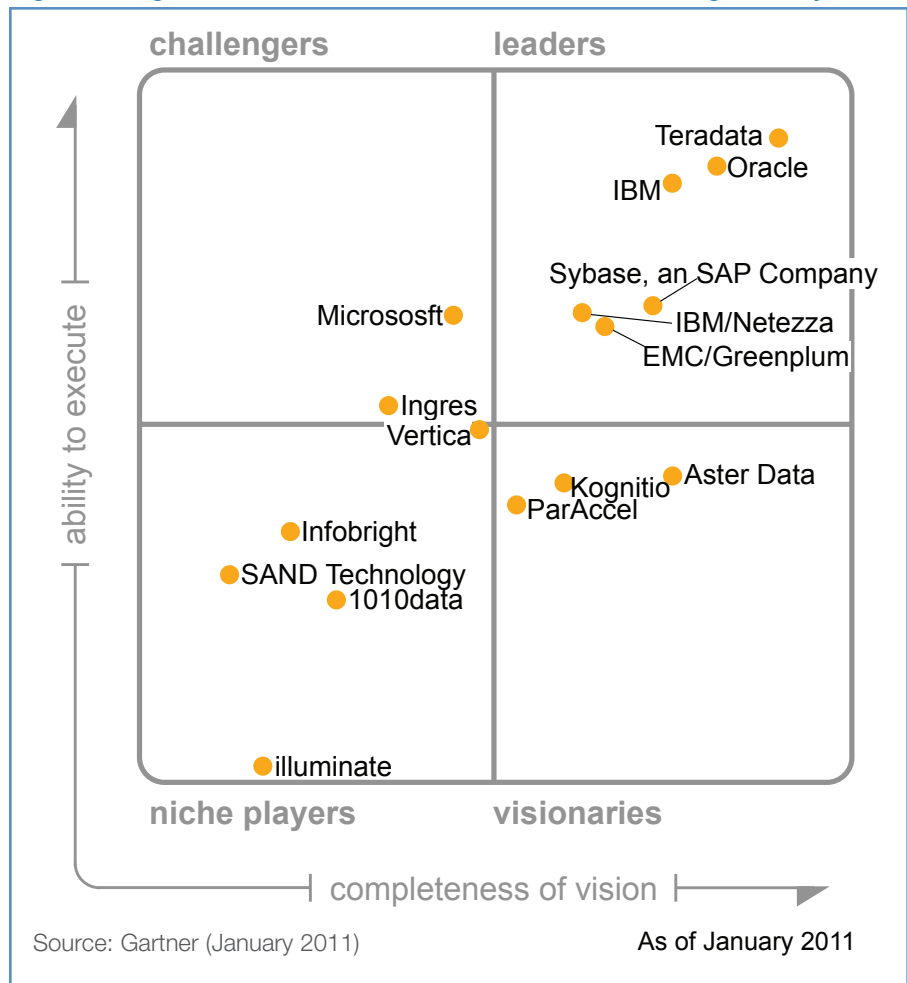
Note 2. Definition of Mixed Workload

The modern complex mixed workload consists of:

- Continuous (near-real-time) data loading — similar to an OLTP workload (due to the updating of indexes and other optimization structures in the data warehouse) — that creates issues for summary and aggregate management to support dashboards and prebuilt reports.

- Batch data loading, which persists as the market matures and starts to realize that not all data is required for "right time" latency, and that some information, being less volatile, does not need to be refreshed as often as more dynamic real-time data elements.

- Large numbers of standard reports — thousands per day — requiring SQL tuning, index creation, new types of storage partitioning and other types of optimization structure in the data warehouse.

- Tactical business analytics in which business process professionals with limited query language experience use prebuilt analytic data objects with aggregated data (prejoins) and designated dimensional drill-downs (summaries). They rely on a BI architect to develop commonly used cubes or tables.

- An increasing number of truly ad hoc query users (data miners) with random, unpredictable uses of data, which implies a lack of ability to tune specifically for these queries.

- The use of analytics and BI-oriented functionality in OLTP applications, creating a highly tactical use of the data warehouse as a source of information for OLTP applications requiring high-performance queries. This is one force driving the requirement for high availability in the data warehouse.

- Fierce marketing and increased resources for winning POCs.

- Continued demand for delivery and management of fully mixed workloads.

- Increased emphasis on supporting new applications with high business value.

- More emphasis on, and appreciation of the value of, the column-store DBMS model.

- Adoption and use of new storage technologies, from disk to solid-state, especially for the management of "hot" and "cold" data.

- With the growing availability and use of storage-class memory, increasing adoption of an in-memory DBMS model.

- More emphasis on, and greater adoption of, cloud-based data warehouse capabilities, albeit for private clouds.

As stated in the previous iteration of this Magic Quadrant, we believe vendors have begun to establish their positions in preparation for a major battle over data warehouse DBMS market share. Vendors that do not differentiate their offerings will either leave the market by choice or be forced out by economic necessity. Once vendors have established their positions, the major fight will begin, probably toward the end of 2013. It is becoming clearer that this will represent a major upheaval in the market, one that the larger vendors need to prepare for and that will give smaller vendors an opportunity. As described in "The State of Data Warehousing, 2011," several aspects of this battle are emerging:

- The combination of repositories, federation and data buses is now possible, given the state of hardware technology.

- The reduced influence of BI platform optimization, in favor of DBMS optimization.

- The increasing influence of master data management and data quality.

- The demand for cloud solutions.

Organizations should increase their emphasis on vendors' financial viability, and closely align their analytics strategies and vendor road maps when choosing vendors.

The data warehouse DBMS market is complex, with a mix of mature and new products. Its complexity reflects many factors, such as:

- The need for DBMS systems to support database sizes ranging from the small to the very large.

- The complexity of data in data warehouses, not only in terms of interrelationships but also of desired data types.

- The fact that data warehouses are built on many different hardware and operating systems, which a DBMS needs to support.

- The growing and regularly changing variety of operations performed in data warehouses, which requires continuous management of the DBMS.

- A DBMS has to support workloads ranging from simple to complex, and to manage mixed workloads in many different combinations.

- The SLAs required by users are shortening, but the implications of not meeting them are more serious.

- The data warehouse has become a "mission-critical system" (see Note 4) in most organizations and therefore requires both high-availability and disaster recovery architectures.

The data warehouse DBMS has evolved from being an information store to a support for reporting and traditional BI platforms, and now into a broader analytics infrastructure that supports operational analytics, performance management, and other new applications and uses such as operational BI and operational technologies (technologies that stream data from devices such as smart meters). Organizations are adding additional workloads with online transaction processing (OLTP) access, and data loading latency is falling to near-continuous loading.

Note 4. Definition of Mission-Critical Systems

Mission-critical systems are systems that support business processes and the generation of revenue, and that, if absent for a period of time determined by the organization and its service-level agreements, must be replaced by manual procedures to prevent loss of revenue or unacceptable increases in business costs. Normally, mission-critical systems require high-availability systems and disaster recovery sites. We include the use of a DBMS as a data warehouse engine in the mission-critical systems category, as we believe that many, if not most, data warehouses in use today fit the definition of mission-critical.

Note 3. Definition of Extreme Data

Issues of "extreme data" arise from the simultaneous and persistent interaction of extreme volume, diversity of data format, velocity of record creation, variable latencies and the complexity of individual data types within formats. "Big data" is another popular term for this concept, but it encourages a focus on a single aspect (volume) and thus creates definitional issues — which will remain unresolved in the market.

There are many other aspects to the data warehouse DBMS market, such as pricing models, geographic reach, partner channels, third-party software partnerships and data warehouse services. We describe these in "The State of Data Warehousing, 2011" and "Data Warehousing Trends for the CIO, 2011-2012."

Market Definition/Description

The supplier side of the data warehouse DBMS market consists of vendors that supply DBMS products providing the database infrastructure of a data warehouse.

For the purposes of this document, a DBMS is defined as a complete software system that supports and manages a logical database or databases in storage. Data warehouse DBMSs are systems that, in addition to supporting the relational data model (extended to support new structures and data types such as materialized views and XML), support data availability to independent front-end application software, and include mechanisms to isolate workload requirements and control various parameters of end-user access within a single instance of the data. This market is specific to DBMSs used as a platform for a data warehouse. It is important to note that a DBMS cannot itself be a data warehouse — rather, a data warehouse (solution/data architecture) is deployed on a DBMS platform. A data warehouse solution architecture can, and often does, use many different data constructs and repositories.

A data warehouse is a database in which two or more disparate data sources are brought together in an integrated, time-variant repository. Its logical design includes the flexibility to introduce additional disparate data without significant modification of its existing entity design.

A data warehouse can be of any size. Gartner defines a small data warehouse as less than 5TB, a medium-sized data warehouse as 5TB to 20TB, and a large data warehouse as greater than 20TB. For the purpose of measuring the size of a data warehouse database, we define data as source-system-extracted data (SSED), excluding all data warehouse design-specific structures (such as indexes, cubes, stars and summary tables). SSED is the actual row/byte count of data extracted from all sources.

Data marts can be deployed on any data warehouse DBMS.

In addition, for the purposes of this document we treat all of a vendor's products as a set. If a vendor markets more than one DBMS that can be used as a data warehouse DBMS, we note that in the section on that vendor but evaluate its products together as a single entity. Further, a DBMS product must be part of a vendor's product set for the majority of the calendar year in question. If a product or vendor is acquired mid-year, it will be labeled appropriately but kept separate on the Magic Quadrant until the following year.

There are many different delivery models, such as stand-alone DBMS software, certified configurations, data warehouse appliances and cloud (public and private) offerings. These are also evaluated together within the analysis of each vendor.

Inclusion and Exclusion Criteria

- Vendors in this market must have DBMS software that has been generally available for at least a year. We use the most recent release of the software for our evaluation. We do not consider beta releases.

- Vendors must have generated revenue from a minimum of 10 verifiable and distinct organizations with data warehouse DBMSs in production.

- Customers in production must have deployed enterprise-scale data warehouses that integrate data from at least two operational source systems for more than one end-user community (such as separate business lines or differing levels of analytics).

- Support for these data warehouse DBMS products must be available from the vendor. We also consider open-source DBMS products from vendors that control or participate in the engineering of DBMSs.

- Data warehouse DBMS or DBMS product vendors that support an integrated front-end tool, but which can also open their DBMSs to competing applications, are included if access is achieved via open-access technology, as opposed to custom-built application programming interfaces.

- Vendors participating in the data warehouse DBMS market must demonstrate their ability to deliver the necessary infrastructure and services to support an enterprise data warehouse.

- Products that include unique file management systems embedded in front-end tools, or that exclusively support an integrated front-end tool, do not qualify for this Magic Quadrant.

Added

None, but see the name changes in the "Dropped" section below.

Dropped

- Greenplum — acquired by EMC on 29 July 2010 — now appears as EMC/Greenplum.

- HP Neoview, as HP no longer actively sells Neoview to new customers.

- Netezza — acquired by IBM on 11 November 2010 — now appears as IBM/Netezza (separately from IBM in this iteration of the Magic Quadrant).

- Sun Microsystems (MySQL) — acquired by Oracle on 27 January 2010 — is no longer a separate vendor. It is included in the analysis of Oracle.

- Sybase — acquired by SAP on 29 July 2010 — now appears as Sybase, an SAP Company.

Evaluation Criteria

Ability to Execute

Ability to Execute is primarily concerned with the ability and maturity of the product and the vendor organization. Criteria under this heading also consider the product's portability, its ability to run and scale in different operating environments (giving the customer a range of options), and the differentiation between data warehouse DBMS solutions and data warehouse appliances. Ability to Execute criteria are critical to customers' satisfaction and success with a product, so customer references are weighted heavily throughout.

Specific Criteria
Product/service includes the technical attributes of the DBMS. We include high availability/disaster recovery, support and management of mixed workloads, speed and scalability of data loading, and support for new hardware and memory models. These attributes are measured across a variety of database sizes and workloads. We also consider the automated management and resources necessary to manage the data warehouse, especially as it scales to accommodate larger and more complex workloads.

Overall viability includes corporate aspects such as the skills of the personnel, financial stability, research and development (R&D) investment, and merger and acquisition activity. It also covers the management's ability to respond to market changes and, therefore, the company's ability to survive market difficulties (crucial for long-term survival).

Under **sales execution/pricing** we examine the price/performance and pricing models of the DBMS, and the ability of the sales force to manage accounts (judging from feedback from our clients). We also consider DBMS software market share.

Market responsiveness and track record covers references (for example, number and size of client companies, nature of configurations and workload mix), general customer perceptions of the vendor and its products, and the diversity of delivery models. We also consider the vendor's ability to adapt to market changes and its history of flexibility when it comes to market dynamics, including use of POCs as required by the market.

Marketing execution explores how well the vendor understands and builds its products in response to the needs of customers (from novices to advanced implementers), and how it develops offerings to meet those needs and the needs of the market in general. We also consider the vendor's geographical ability to deliver solutions.

We evaluate customer support and professional services as part of the **customer experience** criterion, together with input from customer references. Also considered are the track record of POCs, customers' perceptions of the product, and customers' loyalty to the vendor (this reflects their tolerance of its practices and can indicate their degree of satisfaction).

Operations covers the alignment of the vendor's operations, as well as whether, and how, they enhance its ability to deliver. We also include channel partnerships and the vendor's ability to create and use a partnership model.

Completeness of Vision

Completeness of Vision encompasses a vendor's ability to understand the functionality necessary to support the data warehouse workload design, the product strategy to meet market requirements, and the ability to comprehend overall market trends and to influence or lead the market when necessary. A visionary leadership role is necessary for the long-term viability of product and company. A vendor's vision is enhanced by its willingness to extend its influence throughout the market by working with independent, third-party application software vendors that deliver data-warehouse-driven solutions (such as BI). A successful vendor will be able not only to understand the competitive landscape of data warehouses, but also to shape the future of this field.

Specific Criteria
Market understanding covers a vendor's ability to understand and shape the data warehouse DBMS market and show leadership in it. In addition to examining a vendor's core competencies in this market, we consider its awareness of new trends, such as the increasing sophistication of end users, the growth in data volumes and the changing concept of the enterprise data warehouse.

Table 1. Ability to Execute Evaluation Criteria

Evaluation Criteria	Weighting
Product/Service	high
Overall Viability (Business Unit, Financial, Strategy, Organization)	low
Sales Execution/Pricing	low
Market Responsiveness and Track Record	high
Marketing Execution	standard
Customer Experience	high
Operations	low
Source: Gartner (January 2011)	

Table 2. Completeness of Vision Evaluation Criteria

Evaluation Criteria	Weighting
Market Understanding	high
Marketing Strategy	standard
Sales Strategy	standard
Offering (Product) Strategy	high
Business Model	high
Vertical/Industry Strategy	low
Innovation	high
Geographic Strategy	low
Source: Gartner (January 2011)	

Marketing strategy refers to a vendor's marketing messages, product focus, and ability to choose appropriate target markets and third-party software vendor partnerships to enhance the marketability of its products. For example, whether the vendor encourages and supports independent software vendors (ISVs) in its efforts to support the DBMS in native mode.

An important criterion is **sales strategy**. This encompasses all channels and partnerships developed to assist with sales, and is especially important for younger organizations, as it enables them greatly to increase their market presence while maintaining a lower cost of sales. This criterion also includes the vendor's ability to communicate its vision to its field organization and, therefore, to clients and prospective customers.

Offering (product) strategy covers the areas of product portability and packaging. Vendors should demonstrate a diverse strategy that enables customers to choose what they need to build a complete data warehouse solution. We also consider the availability of certified configurations and appliances based on the vendor's DBMS.

Business model covers how a vendor's model of a target market combines with its products and pricing, and whether it can generate profits with this model, judging from its packaging and offerings.

We do not believe that **vertical/industry strategy** is a major focus of the data warehouse DBMS market, but it does affect a vendor's ability to understand its clients. Items such as vertical sales teams and specific vertical data models are considered here.

Innovation is a major criterion when evaluating the vision of data warehouse DBMS vendors in developing new functionality, allocating R&D spending and leading the market in new directions. It also includes a vendor's ability to innovate and develop new functionality in the DBMS, specifically for data warehouses. The use of new storage and hardware models is key. Increasingly, users expect a DBMS to become self-tuning, reducing the resources involved in optimizing the data warehouse, especially as mixed workloads increase. Also addressed here is the maturation of alternative delivery methods such as software-as-a-service (SaaS) and cloud infrastructures.

We evaluate a vendor's worldwide reach and **geographic strategy** by considering its ability to use its own resources in different regions, as well as those of subsidiaries and partners. This criterion includes a vendor's ability to support clients throughout the world, around the clock, in many languages.

Leaders

The Leaders quadrant contains the vendors that demonstrate the greatest support for data warehouses of all sizes, with large numbers of concurrent users and management of mixed data warehousing workloads. These vendors lead in data warehousing by consistently demonstrating customer satisfaction and strong support, as well as longevity in the data warehouse DBMS market, with strong hardware alliances. Hence, Leaders also represent the lowest risk for data warehouse implementations, in relation to, among other things, performance as mixed workloads, database

sizes and complexity increase. Additionally, the market's maturity demands that Leaders maintain a strong vision for the key trends of the past year: mixed-workload management for end-user service-level satisfaction and data volume management.

Challengers

The Challengers quadrant includes stable vendors with strong, established offerings but a relative lack of vision. These vendors have presence in the data warehouse DBMS space, proven products and demonstrable corporate stability. They generally have a highly capable execution model. Ease of implementation, clarity of message and engagement with clients contribute to these vendors' success. Challengers offer a wide variety of data warehousing implementations for different sizes of data warehouse with mixed workloads. Organizations often purchase Challengers' products initially for limited deployments, such as a departmental warehouse or a large data mart, with the intention of later scaling them up to an enterprise-class deployment.

Visionaries

Visionaries take a forward-thinking approach to managing the hardware, software and end-user aspects of a data warehouse. However, they often suffer from a lack of a global, and even strong regional, presence. They normally have smaller market shares than Leaders and Challengers. New entrants with exceptional technology may appear in this quadrant soon after their products become generally available. But, more typically, vendors with unique or exceptional technology appear in this quadrant once their products have been generally available for several quarters. The Visionaries quadrant is often populated by new entrants with new architectures and functionalities that are unproven in the market. To qualify as Visionaries, vendors must demonstrate that they have customers in production, in order to prove the value of their functionality and/or architecture. Our requirements for production customers and general availability for at least a year mean that Visionaries must be more than just startups with a good idea. Frequently, Visionaries will drive other vendors and products in this market toward new concepts and engineering enhancements. In 2010, the Visionaries quadrant was thinly populated with vendors meeting demand from some market segments for aggressive strategies for specific functions, such as the use of MapReduce for large-scale data analytics and massive process scaling in heterogeneous hardware environments.

Niche Players

Niche Players have low market shares or little market appeal. Frequently, a Niche Player provides an exceptional data warehouse DBMS product, but is isolated or limited to a specific end-user community, region or industry. Although the solution itself may not have limitations, adoption is limited. This quadrant contains vendors in several categories: (1) those with data warehouse DBMS products that lack a strong or a large customer base; (2) those with a data warehouse DBMS that lacks the functionality of those of the Leaders; (3) those with new data warehouse DBMS products that lack general customer acceptance or the proven functionality to move beyond niche status. Niche Players typically offer smaller, specialized solutions that are used for specific data warehouse applications, depending on the client's needs.

Vendor Strengths and Cautions

1010data

1010data (www.1010data.com) is a 10-year-old managed service data warehouse provider with an integrated DBMS and BI solution targeted at the business side of organizations, primarily those in the financial and, more recently, the retail/consumer packaged goods (CPG) sectors. 1010data can host its solution using a traditional SaaS model or support a managed solution at the customer's site.

Strengths

- 1010data offers a solution including a DBMS to provide high-speed analytics for businesses. This is a fast-to-market solution — as SaaS — for organizations needing a BI application, lacking BI and data warehousing expertise, or wanting a managed service to complement their internal data warehouse expertise. 1010data's DBMS is fully compliant with SQL and has an Open Database Connectivity interface that can be used for other applications, in addition to its own. Our reference checks and discussions with Gartner clients also show that 1010data is price-competitive with non-SaaS alternatives, especially by reducing the management overhead needed to support a data warehouse environment.

- Since 1010data offers a complete SaaS solution, the customer's business unit and IT organization need little experience of data warehousing or BI. The SaaS model also allows multiple organizations to share large amounts of data without needing to manage it locally — for example, large quantities of CPG data can be shared by multiple retail companies. As a managed service solution vendor, 1010data can complement the customer's internal IT department with fast-to-market solutions for business units, so reducing resource consumption within the IT department. More importantly, the managed service model enables 1010data to leverage software solutions across multiple customers. As new applications are created, they become available to all clients, increasing the availability of those applications to businesses.

- According to our reference checks, 1010 data demonstrates the ability to expand from the financial sector (where it began) into a broader market, including the retail sector. 1010data now claims over 150 customers, and its references support our belief that it is one of the stronger small data warehouse DBMS vendors. In addition, 1010data has seen a growing number of customers install its system on-premises as a managed solution, with several using 1010data as an enterprise data warehouse solution vendor.

Cautions

- With only a fully managed service model, 1010data is susceptible to resistance from IT departments wanting to have all their data warehouses in-house, along with in-house governance of the organization's data assets. To address this issue, 1010data offers to install its system on-premises; however, the system is still managed by 1010data, which raises issues of governance and control for some potential customers. Also, a big challenge for data warehouse SaaS solutions is posed by the issues — real or perceived — surrounding remote locations, security and data transfer performance. And even as these issues are addressed and subside, 1010data will face increasing pressure from cloud DBMS vendors such as Microsoft (with SQL Azure) and salesforce.com (with Database.com).

- 1010data's offering is sold as a fully integrated DBMS and BI solution, which limits potential customers to those wanting a full solution (primarily because of 1010data's pricing model). 1010data's product is a compliant relational DBMS (RDBMS), and customers can use it as a stand-alone system if desired. However, when using 1010data's offering in a stand-alone environment as a data warehouse solution, the cost model is not as advantageous, as 1010data charges the same as if it were managing the offering, and there are internal management costs to add. Customers are advised to check the total cost of ownership (TCO) in such cases, as it may not be advantageous to use 1010data in this way.

- As a solution vendor, 1010data has a different competitive model from vendors of pure-play DBMS offerings. In addition to competing in the data warehouse DBMS market, it competes with system integration vendors that offer outsourced solutions, such as Cognizant and HP (via EDS). Additionally, IBM, Oracle and other large vendors with professional service organizations compete with 1010data in two markets, for both data warehouse DBMSs and services.

Aster Data

Aster Data (www.asterdata.com) sells an MPP DBMS for data warehousing and analytics. Aster Data offers a DBMS, in-DBMS analytics, graph processing, GIS data, blogs, clickstream data, MapReduce applications and more.

Strengths

- Aster Data's nCluster is an MPP DBMS implementation that includes an architecture optimized for in-database procedural processing and analytics. Nodes are assigned specific workload objectives (query management, load/export, parallel processing and backup), which provide for resource balancing during various workloads. Dynamic workload management controls use a rule-based management approach. Aster Data's references report very strong performance with nCluster in all workloads, verifying the capabilities of its dynamic workload management. Aster Data also enables applications, such as analytics written in SQL and/or MapReduce, to execute in parallel on the worker servers. Further, because these applications run in the nCluster product, they are subject to control by the workload manager.

- Aster Data debuted on the Magic Quadrant in 2009 with a strong vision, and in 2010 it added a series of functions and features that indicate continued vision: hybrid row and column store, data store format advisors and hybrid column/row processing (including for MapReduce). Users report that the downloadable Aster Data Developer Express SQL-MapReduce IDE (integrated development environment) simplifies deployment

of MapReduce capabilities. Reference clients report that another strength of Aster Data is its significant scaling capability. Additionally, Aster Data's clients report that with this ease of scaling, they can keep data longer and are rapidly expanding the attributes they capture in data warehouses.

- A recently announced partnership with Cloudera enables the use of MapReduce jobs managed by the Aster Data nCluster DBMS or exported to run on Hadoop Distributed File System (HDFS) files in the Cloudera-managed MapReduce environment. Combined with Aster Data's existing offerings for cloud-enabled deployments on Amazon, Dell, Terremark and AppNexus platforms, this demonstrates Aster Data's ability to work in a heterogeneous analytics environment. Aster Data offers an appliance version of nCluster for Dell hardware, which, when combined with data integration software from Informatica and MicroStrategy or Tableau for BI, enables Aster Data to compete with appliance-only vendors.

Cautions

- Aster Data's customers report that, although its MapReduce functionality exceeds expectations, some of the less sophisticated SQL capabilities are suboptimal (for example, auto-partitioning conflicts with insert/delete, language semantics and other functions). Users report loading issues, and their comments indicate that this is due to table management issues. They also report that patches are not applied correctly or occasionally fail to rectify the intended issues. Finally, they report that the system's performance sometimes becomes unpredictable and recommend enabling Aster Data's dynamic workload management when running heavy workloads. However, Aster Data reports that many of these issues were addressed in nCluster 4.6 (released in September 2010). Aster Data also introduced a new quality assurance process in 2010, which records a significant fall in the number of client-reported issues.

- Aster Data's "mind share" appears limited, as it features in only a small number of competitive situations reported by Gartner clients. This indicates that its market presence needs to grow. We recommend conducting a thorough POC with Aster Data and at least two other vendors; and if MapReduce is to be used, it should be part of the POC. As one of the newer entrants to the data warehouse DBMS market, Aster Data poses a greater risk than the larger vendors.

- Like other small vendors with a solid architecture that differs from the traditional DBMS, Aster Data remains a candidate for acquisition by a vendor wanting to develop, adopt and implement Aster Data's architecture within its own DBMS infrastructure (either as a product in its own right or by reverse-engineering the functionality in an additive fashion). Aster Data will also have to combat the attraction of incumbent vendors noted in the Market Overview section.

EMC/Greenplum

Greenplum (www.greenplum.com) was acquired by EMC on 29 July 2010. Greenplum has an MPP data warehouse DBMS based on open-source DBMS PostgreSQL running on Linux and Unix. It can be sold as an appliance or as a stand-alone DBMS, and has just over 200 customers worldwide.

Strengths

- As we suggested might happen in the previous iteration of this Magic Quadrant ("Magic Quadrant for Data Warehouse Database Management Systems"), Greenplum was acquired in 2010, a development that resulted in the creation of EMC's Data Computing Products Division. This deal moves Greenplum from a small startup DBMS vendor to a position of less risk. It makes available greater funds for R&D, offers stability with EMC backing the technology, and gives the opportunity to leverage EMC's sales and support worldwide. Given the portability of Greenplum's offering — it runs on both Unix and Linux, and has a broad set of hardware options (including those of Dell, Cisco, HP, IBM and Oracle/Sun Microsystems) — the company must now be considered a major vendor in the data warehouse DBMS market. Recently, EMC/Greenplum released its first true data warehouse appliance, the Greenplum Data Computing Appliance, which is sold and serviced by EMC.

- Greenplum has a strong vision and understanding of the data warehousing market. It has demonstrated production scalability to more than hundreds of terabytes. It has also shown the ability to run and manage mixed workloads for a number of references. Through its software architecture, Greenplum can move DBMS code and user-defined functions closer to the storage device, thereby increasing performance. Greenplum supports all major data integration and BI platform vendors, including open-source vendors like Jaspersoft, Pentaho and Talend. In addition, it was the first data warehouse DBMS vendor to deliver a DBMS solution for use in a private cloud infrastructure (Greenplum Chorus). Chorus allows for the creation of a data warehouse environment with self-service provisioning and elastic scale, through the use of a Web portal.

- Greenplum has shown the ability to support many in-DBMS functions, running in parallel, for analytics (open-source analytics, for example), MapReduce, matrix and vector data types, and model building. It is one of the first vendors (along with Aster Data) to support a dual-DBMS model that permits both row-store and column-store tables in the same database. Greenplum was one of the first data warehouse DBMS vendors to implement MapReduce internally for large-scale analytics and to offer external file processing integrated with the DBMS. The latter has enabled Greenplum to manage complex, unstructured data and to connect other implementations of Hadoop MapReduce, such as Cloudera's. In addition, Greenplum now offers a free — though not open-source — single-server version for download and use in development environments running on Mac OS X and Linux.

Cautions

- Now that Greenplum is part of EMC it will find itself competing at a higher level with the mature, incumbent vendors. It must continue to demonstrate differentiation and to support customers accustomed to the type of service provided by a small company. It must minimize the disruption of being acquired by a large company.

- Although strengthened by the EMC acquisition, with fewer than 200 customers EMC/Greenplum remains a relatively small vendor in the data warehouse market, especially compared with the large, mature vendors in the Leaders quadrant. In POCs it finds itself competing with IBM, Oracle and Teradata, but we note that Greenplum does win its share of these. Competition from the traditional vendors, albeit reduced by the acquisition, remains a concern as these vendors have large R&D and marketing budgets and continue to add functionality, which enables them to compete with innovative vendors like Greenplum.

- Greenplum, like other data warehouse-only vendors, will face resistance from prospective customers in situations where it is now possible to use a data warehouse from the incumbent vendor (such as IBM, Microsoft and Oracle). Although EMC is also an incumbent vendor in many organizations, it is an incumbent in other markets, such as storage subsystems.

IBM

IBM (www.ibm.com) offers stand-alone DBMS solutions as well as data warehouse appliances, currently marketed as the IBM Smart Analytics System family. IBM's data warehouse software, InfoSphere Warehouse, is available on Unix, Linux, Windows and z/OS.

Strengths

- IBM caters for most approaches to data warehouse implementation, from custom-built (still the preference of some large IT shops), to preloaded data warehouse appliances, to an appliance-like approach with a fully loaded and configured solution. The wide availability of solutions is the result of IBM's ongoing investment in the data warehouse space. InfoSphere Warehouse, a data warehouse offering based on IBM DB2, is a software-only solution. IBM's data warehouse appliance solution, the IBM Smart Analytics System (formerly IBM InfoSphere Warehouse), is a combined server and storage hardware solution (using the IBM Power Systems server with AIX, or the System x server with Linux or Windows and the IBM InfoSphere Warehouse), complete with service and support. The acquisition of Netezza in late 2010 gives IBM a ready-made Linux-capable data warehouse appliance, which competes directly with Oracle's Exadata.

- IBM's introduction of InfoSphere BigInsights reflects its strategy for adapting/adopting the open-source Hadoop project, and includes offerings to aid the design, installation, integration and monitoring of the use of these open-source technologies within an IBM-supported environment. By tying together relational data, data streams and Hadoop files, IBM's stack builds confidence among managers of existing warehouse implementations that the product is evolving as new demands emerge.

- IBM is the only DBMS vendor that can offer an information architecture (the Information Agenda) across an entire organization, covering information on all systems, including OLTP, data warehousing and retirement of data (with its Optim products). In addition, Optim Database Administrator can propagate schema changes from test to production environments. Another performance optimization feature is partitioned updates to cubes for real-time analytics. This is very compelling for organizations in which IBM is the incumbent vendor, and IBM is good at using the Information Agenda for data warehousing. IBM maintains strong support from its very large customer base.

Cautions

- From our Magic Quadrant survey for 2010, it was clear that IBM's customers still detect a shortage of skilled implementers such as architects and DBAs. However, this growing concern of customers is also an encouragement for vendors in that, with demand exceeding supply, it indicates that the market is succeeding. Clients also report that IBM's support appears disconnected from its product strategy, in that support seems incident-focused, rather than focused on general solutions and practices. Clients report that increases in their internal staff numbers are driven primarily by the need for skilled architects, modelers and DBAs, which indicates that users are trying to solve support issues themselves. Gartner observes that the complexity and volume of data under management increased dramatically in 2010, which is also driving up demand for skills — and again indicates that market success is increasing demand for expertise.

- In 2010, Gartner clients reported that IBM was selected 85% of the time, when IBM was a candidate. Normally, a high win-rate is a strength, but in this case there is a mixture of cautious optimism and valid concern. IBM has embarked on a mission to qualify its prospects better for warehousing, and is therefore competing for fewer, better-qualified prospects — a solid tactical decision that could nevertheless jeopardize its execution. However, even with enhanced qualification, 27% of current IBM customers selected another vendor when choosing a warehouse platform. This means that at least 27% of IBM's current DBMS customers are willing to deploy a different DBMS for the warehouse, and the percentage is probably at least double that. This, in turn, means that IBM is at least passively refusing to compete for business from some of its existing customers. On the positive side, such decisions indicate that IBM is aware of its product capabilities and delivery capabilities, and is not trying to sell products or commit resources to poorly matched opportunities.

- In the last iteration of this Magic Quadrant, Gartner indicated that IBM needed to grow at least at the same rate as the market. In 2009, the RDBMS market was flat, but IBM's market share declined by about 0.7 percentage points. This was, however, better than Oracle's decline of 1.8 percentage points, and we believe that, as the market recovers from the economic crisis, IBM will return to growth at a rate faster than the market average.

IBM/Netezza

Netezza (www.netezza.com) was acquired by IBM in late 2010. It markets its TwinFin platform, which is based on IBM's System x platform, continues to pursue a hardware acceleration strategy with multilayered processing, and has introduced complex and large dataset processing beyond the warehouse. Specifically, its work with ISV partners capitalizes on its architecture's use of processors.

Strengths

• Netezza continues to mature its product by following a customer-driven road map, with the major addition of a built-in capability, called i-Class, that brings in-DBMS, parallel analytics to the Netezza engine. Netezza has also continued to develop partnerships with many software vendors to extend and support its platform — as shown, for example, by the Netezza Data Virtualizer powered by Composite Software. In partnership with EnterpriseDB, Netezza has added an Oracle compatibility front-end to TwinFin, enabling Oracle PL/SQL and Oracle SQL to execute natively on its platform. Additionally, Netezza continues to evolve its product with additional system administration, workload management and data management enhancements, pushing it toward becoming a complete data warehouse platform. Besides the in-DBMS analytics, specific features added during the past few years include recovery from S-blade failures, data compression, auto-regeneration of disk-stored data after a failure, and system and query statistical metadata for active optimization.

• Netezza's introduction of TwinFin brought a much-needed physical separation of its multiple levels of processing technology. The move to a standard hardware architecture (using IBM's System x BladeCenter, along with Netezza's proprietary field-programmable gate array [FPGA] technology) enabled Netezza to market a modular, upgradable and scalable appliance. This may have led to the acquisition by IBM. The move to standard hardware architecture, such as x86, has also increased the potential for third-party software partners. Netezza continues to form partnerships with vendors wishing to run their application code on its product's processors. Thanks to Netezza's architecture, the effort involved in doing this is relatively small, and the result is very impressive, with increased parallelism and better performance from applications. In addition, in 2010 Netezza added NEC as a partner: a Netezza appliance is being built on NEC's x86 platform. It is intended for the Asian market, and will be sold and supported by NEC.

• Judging from customer references and discussions with Gartner clients, Netezza's solution is relatively simple to install, implement and maintain; they report some of the shortest time frames in the market from delivery to production. In addition, references report that Netezza continues to improve its complex workload management capabilities. The company has dispelled the perception that it is a data-mart-only appliance vendor. Further, conversations with Gartner clients indicate that Netezza continues to compete well on both price and performance against Oracle's Exadata. That Netezza had over 500 customers at the end of 2010 points to the same conclusion.

Cautions

• Netezza held its own against the megavendors that entered the appliance space in 2009 and 2010, and the acquisition by IBM will boost its competitive positioning. However, the extent of this boost will depend heavily on TwinFin's positioning by IBM in relation to IBM's Smart Analytics. Although the acquisition could help Netezza win business from customers for whom IBM is the incumbent vendor, it could have the opposite effect where IBM is not the incumbent. Further, the positioning within IBM's data warehouse product portfolio will be critical. It will require clear communication of marketing messages to IBM's field marketing and sales forces, and to current and prospective customers, to prevent confusion over the data warehouse offerings of Netezza and IBM.

• Netezza is very good at isolating POC constraints when competing head-to-head with other vendors. Prospective customers are advised that POC results, while excellent and valid, are often based on isolated workload situations (of the single-workload type) or leverage Netezza's massive hardware strategy. With TwinFin, this is reported to be less of an issue, but we still advise customers to do complex workload testing as part of any POC.

• Netezza's prices are no longer a disruptive force that give it a competitive edge. Other vendors have responded with similar prices and discounts, and have introduced entry-level solutions, such as Oracle with Exadata and Teradata with its 2650 appliance. Prospective customers should no longer assume that Netezza wins automatically on price. However, TwinFin, with its new architecture based on standard IBM System x hardware, will help Netezza here. Furthermore, Netezza has been successful at moving the discussion away from low prices to one of price/performance, which Gartner believes is more important than low prices alone. On the other hand, Netezza's long-established differentiation in terms of CPU per disk drive ratio is diminishing as other vendors reduce theirs: Oracle now has one core per disk drive in its Exadata Storage server, while the Teradata 2650 has one core for every two drives, with one thread per drive.

illuminate

A small software vendor, illuminate (www.illuminateinc.com) has an integrated data warehouse DBMS (iCorrelate) and BI tools. The focus of the system is to store all potential relationships between any data element in the database and any other data element. The company has just under 100 customers. They are mainly located in Spain and other parts of Europe, with a few in the U.S. and Latin America.

Strengths

• Customers continue to report that storage capacity is small and query performance fast. The solution stores abstracted data values as a metadata master set in the database. This, along with some tools from illuminate, enhances data quality by ensuring single storage of each value. A purely column-vectored approach reduces the volume of the database, as repeated values within a column are addressed, although repeated values are still possible in the overall database. The solution's correlation theory and use of metadata eliminates those remaining multiuse redundancies.

- The proprietary technology, which is simple to install and use, is shielded by traditionally understood query language and system-level semantics. This structure is automatically built and maintained by illuminate's DBMS intellectual property as data is loaded, so DBAs used to row- or column-vectored, hierarchical files, or any other data file management system, do not have to develop custom-load processing.

- Query processing is enhanced, as the solution effectively creates prejoins for all existing data relationships in the data model. The process is repeated when new datasets are added, with an effect almost like a spider's web that stores every correlation that can be inherited from the data already held in the database, as well as from any newly added data.

Cautions

- In 2010, illuminate almost disappeared from Gartner client inquiries. This does not mean it is not winning new customers but, without any significant capability to penetrate new markets or segments, its viability is at risk. On the positive side, the organization has a very low "run rate" for expenditure and is therefore more easily supported by a small base of customers and investors. It has had some success in North America since opening its first office in the U.S. over three years ago. Lack of market presence jeopardizes illuminate's technical capabilities, but a new CEO and changes to its marketing approach show some promise.

- The company's partner program for distributors and vendors of third-party software is ineffective. Gartner noted in 2008 that illuminate would have to leverage partners and channels if it wanted to gain mind share in this market. In 2009, illuminate expanded its partner channel to about 48 partners and two OEMs worldwide, but although initially successful in 2010, this approach seems to have stalled, judging from Gartner inquiries. On the other hand, illuminate did recently secure a small number of new customers in North America.

- illuminate has been unable to articulate the advantages of its unique technology to the market. This would put it at greater risk were its financial viability to come into question. Although illuminate's technology has specific advantages, it does not appear to be a key acquisition target, probably because of its complex delivery model and internal architecture.

Infobright

Infobright (www.infobright.com), which has offices in Canada, Europe and the U.S., offers a combination of a column-vectored DBMS and a fully compressed DBMS. The company offers both an open-source version (Infobright Community Edition [ICE]) and a commercial version (Infobright Enterprise Edition [IEE]).

Strengths

- Infobright offers the only open-source column-store DBMS on the market. This puts it a unique position, with over 120 IEE customers at the end of 2010. Since it began offering ICE and IEE in 2008, it has matured and introduced commercially

licensed and open-source licensed offerings. There is considerable differentiation between the two products, with IEE including additional features for performance, warrantee indemnification and services. Infobright integrates MySQL's interfaces with the DBMS, so that customers can leverage existing tools (both data integration tools, including the MySQL loader, and BI tools). This enables Infobright to replace MySQL DBMS infrastructure more easily. Because Infobright has an open-source pricing model for ICE (no license fees) and a low-cost model for IEE (based on the amount of SSED), its cost model makes it very interesting for organizations wanting to optimize data warehouse costs.

- The Knowledge Grid in-memory metadata store is a major differentiator for Infobright as this analyzes queries to minimize the number of "data packs" that have to be decompressed to give a result (data packs are the compressed domains/regions of data in Infobright's offering). Decompressing data in memory is already faster than reading full-volume data on disk, so this further enhances performance by limiting decompression to the data needed. Also, the Knowledge Grid sits above the data packs, adding an additional set of metadata and enabling even greater performance, according to Gartner's reference checks.

- Infobright is now focusing on machine-generated (operational technology) data, normally streaming data that is stored and then analyzed. This makes it a good choice for Web-generated data (such as clickstream data) and telco data. Infobright references report very fast analysis of this type of data, primarily due to the high degree of compression and the Knowledge Grid. Infobright is also gaining traction as a DBMS OEM for software vendors that analyze this type of data. Additionally, Infobright is increasing its partnerships with other software vendors, and is part of an open-source reference architecture alliance for BI and data warehousing that includes Pentaho, Jaspersoft and Talend.

Cautions

- As a small, relatively young vendor, Infobright must continue to differentiate its offerings and open-source model from mature column-store DBMSs. It will be challenged by more established vendors as they begin to offer column-store and column-compression alternatives. As pure analytic workloads are becoming more complex, Infobright must move quickly to manage these workloads — something its products cannot yet do — even if it desires to remain an analytic DBMS engine vendor.

- Infobright has been slow to achieve revenue growth from its commercial product, and to demonstrate that the effect of its open-source product is not to reduce its revenue. If the open-source version is good enough, customers may opt for it, rather than the commercial version, though Infobright may be able to license some of its technology (as EnterpriseDB has done) to increase revenue. Its distinct technology and low revenue also make it a likely acquisition target, especially for its Knowledge Grid and column-store technology.

- Infobright makes extensive use of portions of MySQL using the OEM version of MySQL under the General Public License. Although Infobright has a long-term (five-year) contract with Sun for MySQL, risks remain due to the uncertain future of MySQL following the Oracle acquisition. We believe that Infobright can add compatibility with other DBMSs, such as PostgreSQL, which would help stabilize its future, reduce some customer concerns, and add flexibility to its products.

Ingres

Ingres's (www.ingres.com) solution is a general-purpose DBMS with a 30-year history as one of the original RDBMS engines. Its solution is now open-source. The company has many customers running mission-critical applications, including data warehouses.

Strengths

- Ingres, a mature vendor, has more than 10,000 customers using its DBMS, and our reference checks show them to be very loyal. Most have OLTP applications, but Ingres also has its share of smaller data warehouses (up to about 2TB). Ingres has converted almost all its pre-open-source customers to open-source subscriptions. Ingres's is the only open-source DBMS with a substantial number of data warehouse customers, especially for database sizes greater than several hundred gigabytes. In addition, it is the only open-source DBMS with proven maturity for mission-critical applications, including data warehousing. With VectorWise technology and strong software partners, Ingres is positioned well as a Challenger.

- Ingres has gained many third-party software partners, specifically in the BI market. An example is the open-source BI vendor Jaspersoft, which offers a software appliance (or bundle) with Ingres for BI. This is the primary driver of new installations in data warehousing, with both new and existing customers looking for an open-source stack that supports BI.

- Ingres contains most of the features necessary for data warehousing, such as partitioning, compression, parallel querying and multidimensional structures. In addition, in 2010 Ingres released the VectorWise engine based on research done in the Netherlands to make use of the internal instruction parallelism and cache of the x86 processor. This greatly increases the performance of Ingres, especially in analytic applications. With new server platforms emerging with storage-class memory (of 1TB and more), VectorWise will prove a valuable asset for data warehousing and analytics as more of the data warehouse moves to memory.

Cautions

- Although, with VectorWise, Ingres recently enhanced its ability to support analytic data marts, the company must address the areas of enhanced data warehouse functionality, storage management and mixed workload management if it is to compete with larger, more mature vendors in the data warehouse DBMS market and meet the needs of the broader data warehouse market.

- Although Ingres offers professional services in data warehousing and has a go-to-market strategy with its partners, it lacks data models and the necessary marketing and sales expertise for data warehousing. Also, although Ingres has the strongest open-source DBMS offering for data warehousing, the open-source marketing model for data warehousing is weak.

- Ingres's 30-year history works against it, given that it has not regained much market traction. This is an issue of market perception, which is difficult to change. Although Ingres has gained new customers and third-party relationships since becoming an open-source company, to become a serious competitor in this market it must continue to show increased growth in both revenue and numbers of new customers.

Kognitio

Kognitio (www.kognitio.com) started by offering data warehouse appliances and warehousing as a hosted service. Today, it has a mixture of customers using its DBMS (WX2) separately as an appliance, a data warehouse DBMS engine, or data warehousing as a managed service (hosted on hardware located at Kognitio's sites or those of its partners).

Strengths

- In 2010, Kognitio added multilingual support, and reference customers reported significant concurrency capabilities. WX2 version 7 already included in-memory analytics, and references continue to report that the speed of query and load performance is excellent. In addition, the DBMS is already an in-memory DBMS, with hot data held in-memory and cold data on disk, as managed automatically by the DBMS. A data warehousing as a service (DaaS) model permits clients to expand their warehouse incrementally, and clients note that this model provides for low upfront costs with virtually no capital expenditure required to get started. Customers also report excellent support and product management.

- Kognitio pioneered the DaaS model by which a data warehouse DBMS is delivered as a managed service from the DBMS vendor. Clients buy data warehousing services from Kognitio, while Kognitio hosts the database. This is a growing segment of the data warehouse DBMS market. Kognitio's customers report that deployment of large-scale data warehouse efforts takes as little as 10 weeks using this model. Kognitio also works with deployment partners such as Capgemini (Kognitio contributes to Capgemini's Immediate cloud computing offering). In addition, and in line with market demands, Kognitio has an appliance to install on-site for customers wanting their own infrastructure.

- Kognitio opened offices in the U.S. three years ago and is developing partnerships to sell its product. This has started to produce results, with several new customers coming from these partnerships. Kognitio has also added several hosting partners in the U.S. and the U.K. that offer managed services on WX2. The U.S. presence and additional partners have enabled Kognitio to grow, despite the sluggish recovery from the economic downturn.

Cautions

- Customers indicate that the lack of high-availability options is bothersome. The vendor reports that version(s) 7.x include features such as automatic disk space reclamation and automatic recovery from node failure, but Gartner has not received comments from references on these features. In addition, interoperability with popular BI tools, such as those of IBM (Cognos) and SAP (BusinessObjects), is difficult to manage, and not all features of the BI platforms are supported. This problem is compounded by Kognitio's small market penetration and the resulting scarcity of tool expertise in the market.

- Kognitio has a very substantial opportunity in the small or midsize business (SMB) data warehouse and BI market. However, the SMBs' "darling" worldwide is usually Microsoft, and other major DBMS vendors also offer managed data services. Kognitio will need to work more closely with professional services partners to protect itself against the megavendors that now come armed to every competitive bid with their own professional services, appliances, managed data services, models and methodologies for delivery. Another strategy would be to try to become the No. 1 brand in specific industries or geographic markets.

- Kognitio remains a small vendor with fewer than 50 customers worldwide. This makes it increasingly difficult to sell to organizations that have incumbent vendors, and to compete with some of the lower-priced appliance offerings.

Microsoft

Microsoft (www.microsoft.com) continues to market its SQL Server 2008 DBMS for data warehousing customers that do not require an MPP DBMS. Microsoft released its own MPP data warehouse appliance, the SQL Server 2008 R2 PDW, in November 2010, but the date of its availability did not allow us to consider it when deciding Microsoft's position in the present Magic Quadrant.

Strengths

- Microsoft continues to offer value for the price customers pay, giving high value with a low TCO. Buyers of SQL Server 2008/R2 Enterprise Edition also receive SQL Server Analysis Services (SSAS), SQL Server Reporting Services (SSRS) and SQL Server Integration Services (SSIS), Master Data Services and StreamInsight, which means that online analytical processing (OLAP), reporting and data integration for extraction, transformation and loading (ETL), master data management and streaming data processing are included in the entry price, although these capabilities are normally deployed using separate servers. The license price — currently listed as $27,499 per socket for the Enterprise Edition — is also lower than that of many other vendors that price by CPU or core.

- The much anticipated release of the PDW occurred on 8 November 2010. With it, Microsoft also launched a new support and services offering for Tier 1 customers. We believe the PDW resolves some of the scaling issues with SQL Server as a warehouse DBMS — but clustering remains a more manual process with Microsoft's solution than with some of its competitors' offerings. SQL Server as a DBMS has a large installed base (accounting for nearly 50% of Windows DBMS revenue), and we believe that small and midsize data warehouses form a large part of this base. In addition, Microsoft offers SQL Server Fast Track Data Warehouse, which includes validated reference architectures for building a balanced data warehouse infrastructure.

- According to our reference checks and discussions with our clients, worldwide support from Microsoft is extensive, encompassing partners, value-added resellers, vendors of third-party software and tools, and widely available SQL Server skills. This broad support, partner and channel strategy serves Microsoft well for delivery and execution in the data warehouse market, and makes it an example to other vendors.

Cautions

- Gartner's interactions with clients and reference checks find clients reporting issues with clustered server performance, difficulty with high availability using active-passive server clustering, and a relative lack of performance-monitoring tools specifically related to SSIS. Clients also report that these issues lengthen implementation cycles and create daily operational difficulties. We believe that Microsoft will address many of these issues in future releases of the DBMS (for example, in the SQL Server code named Denali) and with the PDW, but they inhibited adoption in 2010. The lack of attention to high availability, clustering and management, coupled with a late-to-market MPP solution (PDW), shows that Microsoft has generally not understood the market's direction and needs before other vendors.

- Much of what PDW delivers is already offered by data warehouse appliances from other vendors, but, although late to the market, PDW is just in time for an entire wave of late adopters of data warehouse and BI strategies. Another issue is that, since the full SQL Server offering comes with PDW, including SSAS, SSRS and SSIS, the temptation will be to scale the entire PDW environment as a single entity, whereas other appliances have shown that data integration, data warehouses and BI optimization scale independently of each other.

- SQL Server runs only on Windows Server and therefore lacks the portability of most of its competitors. Although Microsoft considers this an advantage (due to tighter integration of SQL Server with the operating system), some IT organizations do not consider SQL Server an option, as they are not willing to run production DBMS infrastructure on Windows Server in a data center environment. However, in the appliance market, PDW provides a comprehensive solution and portability becomes less important.

Oracle

Oracle (www.oracle.com) remains a leader in data warehousing, with 48% of the RDBMS market. In 2008, Oracle introduced its first data warehouse appliance, the HP Oracle Database Machine (Exadata). In 2009, Oracle changed the platform from HP to Sun Microsystems (Exadata V2), and in 2010 announced its third generation of Exadata. It now offers the Oracle Exadata Database Machine X2-2 with eight 12-core nodes, and the Oracle Exadata Database Machine X2-8 with two 64-core nodes. Also in 2010, Oracle completed its acquisition of Sun, including the MySQL open-source DBMS.

Strengths

- Oracle gives customers a wide variety of choices, including what Gartner believes to be the most portable DBMS. Oracle has three distinct data warehouse solutions: Oracle Database 11g (the stand-alone DBMS); Oracle Reference Configurations (certified server and storage configurations); and Exadata (X2-2 and X2-8), now Oracle's recommended data warehouse platform — which includes a DBMS appliance (the Oracle Exadata Storage Server) with storage optimized for data warehouses based on Oracle Database 11g Release 2 (11gR2), Oracle Real Application Clusters (RAC), Automatic Storage Management (ASM) and Sun x86 hardware (sold and serviced by Oracle). Oracle continues to extend the stack to hardware, meaning customers have a single vendor for support. Oracle is increasingly gaining acceptance in the market with Exadata, just two years since its release, judging from the number of inquiries we receive about it.

- Oracle Database 11g has added enhanced materialized view and cube management (notably, transparent SQL access and incremental update capability). This increases Oracle's ability to deploy end-user optimization layers with features not found in other DBMSs. Oracle Database 11g also brings enhancements to Oracle's partitioning option; these include Partition Advisor, which suggests types of partitioning to enhance performance based on the database schema. Finally, with Exadata, Oracle has Smart Scan (to offload some DBMS functionality to the storage server), Hybrid Columnar Compression (which reduces storage requirements and increases performance), and Exadata Smart Flash Cache (up to about 5TB of flash memory to optimize data access and queries), which gives Exadata up to a tenfold increase in average performance compared with Oracle's performance on stand-alone hardware (according to Gartner's clients and Oracle's references).

- Oracle RAC with ASM (available in a stand-alone DBMS and on Oracle Exadata) is widely accepted as an enterprise-level DBMS platform for data warehousing, capable of supporting large data warehouses (defined in the Market Definition section as bigger than 20TB) — see "Oracle RAC Moved to Mainstream Use." The scale-out configuration allows for flexibility (adding servers and storage without downtime), while providing a base for the high availability required by the new data warehouse SLAs that are being implemented. Oracle RAC is also the primary software component enabling multiple, connected (using InfiniBand) Exadata machines to operate in a cluster for implementations requiring more processors and/or storage.

Cautions

- From discussions with Gartner clients desiring POCs, we have learned that Oracle prefers not to perform on-site POCs. Instead, Oracle has opened nine international test sites for Exadata, and is pushing customers to use these, instead of on-site installations, for POCs. Gartner always recommends a POC to prove that a platform can handle the required workload. We also recommend on-site POCs for all data warehouse infrastructure sourcing decisions, and that clients press Oracle to perform on-site POCs with the client in control of all aspects of each POC.

- Although with Exadata, Oracle has reduced the full-time equivalent (FTE) staff requirements for the administration and optimization of the data warehouse, Gartner's clients continue to report that Oracle's FTE requirements are higher than those of some other data warehouse DBMS vendors, primarily in stand-alone environments. This, however, has become less of an issue during the past year, due not only to Exadata, but also to many of the new functions and tools available with Oracle Database 11gR2, such as the Automatic Database Diagnostic Monitor now available for Oracle RAC, the new partitioning advisor tool and the Oracle Database Resource Manager. However, it should be noted that many of the DBA productivity tools are optional products with additional license costs.

- Gartner clients continue to identify Oracle's pricing and contract practices — for example, its high prices (for some configurations), uneven and wide-ranging discounts, increasing software audits, high cost of maintenance and reluctance to negotiate on renewals — as greater issues than with other vendors. Recently, this has become a standard topic of inquiries from Gartner clients who fear being "locked in" to Exadata if they purchase the entire software stack from Oracle. For Oracle appliances to become widely accepted, the company must address these issues by introducing standardized discounts, flexibility in the maintenance model (for example, allowing for reductions in maintenance fees as product licensees are no longer required) and removing contract terms specifically designed to lock the client further into Oracle (such as those insisting that clients pay maintenance costs even for unused licenses).

ParAccel

ParAccel (www.paraccel.com) makes its second appearance on the Magic Quadrant. This company's software solution includes the ParAccel column-vectored database and storage management interfacing/management.

Strengths

- ParAccel has many multibillion-dollar enterprises among its clients. They gain specific technological advantages from its solution when performing analytics (see below), partly due to a price/node model that scales with the performance requirements of the end-user organization (and has a per-socket price as opposed to a per-core price). In addition, ParAccel can gain clients through partnerships with vendors of BI and data integration tools, and offer complete solutions with those partners.

- ParAccel easily combines disk utilization with memory utilization in query processing. Its customers routinely join tables containing millions of records during query execution. These include self-joins in analytics such as market basket analysis and drug interaction analysis, as well as support for asset-backed securities analysis, credit-card fraud analysis and risk exposure analysis (regarding liquidity, for example). ParAccel performs well in many POCs, with references reporting POC tests against many of the high-performance vendors such as IBM, IBM/Netezza, Oracle and Teradata.

- With about 40 customers in the pharmaceutical, retail, financial and media/advertising analytics sectors, ParAccel has a good reference base. This base indicates significant performance capabilities, as compared with incumbent solutions and shown by competitive bid processes that include POC evaluations. References specifically note that fully normalized data does not require an optimization layer (marts, materialized views and cubes) to achieve superior performance.

Cautions

- References report software quality issues in relation to upgrades. Indications are that the "leader" node creates a single point of failure, and frequent restarts are often the result during final production implementation and upgrades. ParAccel reports that all nodes have a "hot standby," but this conflicts with clients' reported experiences. It is unclear whether the hot standby is simply not being used by these clients or whether there are problems with it.

- ParAccel offers services and customer support commensurate with its size, but must be ready to scale up its customer support to tackle the new issues that will inevitably arise for any successful startup. Early indications from reference customers show that ParAccel is meeting support needs at their current rate of growth, but we note that their growth in 2010 was commensurate with the company's size and are concerned that, if such growth continues, ParAccel's support model could come under strain — a good problem to have, but a problem nonetheless.

- ParAccel is a recent entrant in a very big market with many vendors. This means that, as with any new entrant, either client organizations must be willing to augment their corporate standard analytics technology/vendor with ParAccel, or ParAccel must displace the corporate standard by overtaking market share leaders in terms of mind share and then beating them on price and in POCs. Indications in 2010 were that the POC model demonstrates ParAccel's advantages, but that becoming the corporate standard for analytics remains a challenge. ParAccel reports that it is using departmental and limited use cases to generate repeat sales to existing customers that are expanding their implementations. In the short term, ParAccel is likely to continue to compete in the specialty category of recursive, very large data analysis for departmental users.

SAND Technology

SAND Technology (www.sand.com) is a column-store DBMS vendor. It has been in existence for approximately eight years and has fewer than 100 customers. SAND uses techniques such as tokenization and compression to strengthen its column-store design. Its technology is used as an analytic engine and as an archive engine.

Strengths

- In 2010, SAND altered its market positioning to present itself as a data analytics platform. To this end, it introduced text search capabilities to its column-store DBMS (sound/spell like, relevance ranking and other text-based capabilities), as well as cloud support functionality (shared processor/storage and distributed processing management). The addition of managed, dependent, disconnected data marts enables synchronization and updates to intermittently connected data marts. As an archive tool, SAND's solution achieves greater compression than other DBMSs because of its use of tokenization in addition to the column-store, and the resulting archive is SQL-accessible.

- Almost all reference customers report that the compression rate of SAND's column-store DBMS is impressive. Additionally, those using it as an archive or an enhancement to SAP's Business Warehouse Accelerator report solid integration, although direct interfacing proves more difficult when it is the primary warehouse. SAND refers to its core engineering as "infinite optimization," and because of the tokenization and column store, it requires no indexing or query tuning. It is also a good choice for analytic data marts to support the off-loading of workloads from an enterprise data warehouse. In addition, several customers use SAND's technology as an enterprise data warehouse.

- SAND's client base remains loyal. With new clients being slowly added from its partnerships with Accenture, Open Text, SAP and TG-Energy, and from the attraction of native SAND products, it could remain a viable vendor in this market or be acquired for its technology — either would be good for its customers. SAND is pursuing cloud business with a DBMS deployed on mobile devices with support for synchronization — a viable strategy given the small size of the stored datasets.

Cautions

- Because of SAND's smallness — it has fewer than 100 customers — it will continue to struggle against the larger vendors and venture-funded startups that can invest more in R&D, marketing and sales. This is an issue for all column-store DBMS vendors.

- SAP's acquisition of Sybase poses a technological challenge for SAND. Sybase IQ is also a column-store database, and now that SAP has its own technology, its partnership with SAND will probably wither. Additionally, SAP has Sybase Adaptive Server Enterprise (ASE) and its own in-memory technology

to accelerate performance. Nor does the fact that Sybase also provides mobile data management solutions help SAND. However, SAND anticipated this development and increased its focus on direct channels in 2010, as the partnership with SAP has continually produced direct customers for SAND. Then again, the potential loss of SAP's powerful marketing channel is a concern.

- Customers report primarily niche use of SAND's column-store DBMS, to support established warehouses. As other DBMS competitors develop their own archiving and information life cycle management strategies, it will become more difficult for SAND to maintain this revenue stream. SAND must gain mind share for its new role as an analytics platform vendor.

Sybase, an SAP Company

In 2010, Sybase (www.sybase.com) was acquired by SAP. Although Sybase has several DBMS products, our analysis is based on Sybase IQ, which was the first column-store DBMS and is Sybase's primary data warehouse DBMS. It is available as both a stand-alone DBMS and a data warehouse appliance, through several system integration vendors. In 2010, Sybase released the first version of Sybase IQ to support an MPP architecture.

Strengths

- During the past few years, Sybase has shown increased ability to move from offering an analytic data mart to offering an enterprise data warehouse DBMS. It has added substantial mixed workload management, faster loading capabilities (to address the biggest issue with column-store DBMSs), query parallelism across multiple processors, and now, with Sybase IQ 15.3, the ability to scale horizontally across a cluster of servers with MPP capabilities. Additionally, Sybase has added features to IQ such as integrated text search and analysis, in-database data mining, and Web-enabled language drivers such as Python, PHP and PERL — each targeted at a new generation of analytical applications. Recently, we have learnt from users of our inquiry service that Sybase IQ is being considered and selected as a complete data warehouse solution. The company's real-time analytics solution, Sybase RAP — The Trading Edition, which includes Sybase CEP for complex-event processing (CEP) and a built-in package for time series analytics to support demand for CEP, has seen solid adoption in the financial services sector since its introduction in 2009. RAP is also available as a general real-time analytics platform for CEP. In January 2010, Sybase acquired the Aleri Streaming Platform to help it build CEP applications for RAP.

- Sybase IQ achieves data compression ranging from two to 10 times, depending on the data's structure. Because analytics typically uses fewer columns but larger numbers of rows, Sybase IQ performs very well for analytic applications. The company has consistently won POCs with analytic applications, sometimes with performance 100 times greater than its competitors. This makes Sybase IQ an extremely desirable DBMS platform for an analytic data mart, to optimize and enhance an organization's overall data warehouse architecture.

- With its acquisition by SAP, Sybase has gained a stronger position in the market. SAP brings a larger sales force, a strong application platform, increased funds for R&D and general financial stability. With Sybase now certified as a DBMS platform for SAP applications, we believe Sybase will achieve increased market share growth in 2011 and beyond. Judging from Gartner inquiries, the main inhibitors of Sybase adoption in the past have been a low market share and the perception that Sybase is too small to remain viable — but following the acquisition by SAP, these no longer apply.

Cautions

- Although Sybase IQ has a large installed base, with over 2,000 customers, it faces competition from data warehouse DBMS vendors, such as Aster Data and EMC/Greenplum, that have introduced column-store capabilities, and others, such as Oracle, with column-based compression within row-vectored DBMSs. Note that these other column-store models are not yet complete column-store integrated systems like those of Sybase and Vertica. However, we believe that during the next 12 to 24 months column-store DBMS features will become more pervasive in row-store DBMS engines as an alternative storage model, and that this will pose a threat to all column-store-only DBMS engines.

- As Sybase continues to expand into the enterprise data warehouse space, it will face increased competition from incumbent vendors, and POCs will become more difficult. Although Sybase IQ remains ahead of the column-based newcomers in execution and has shown the ability to scale to data warehouse solutions, its challenge will be to continue to respond to new market demands by offering a wider variety of data warehouse solutions and moving customers on to a full-scale data warehouse solution.

- Sybase has enjoyed a strong relationship with IBM's Power Systems division, resulting in the Sybase Analytic Appliance, which is sold and supported by third-party system integrators. With the acquisition by SAP, and IBM's acquisition of Netezza and stronger efforts to sell IBM products, we believe that IBM's interest in the Sybase Analytic Appliance will diminish, which will lead to it being discontinued. On the other hand, in 2010 Sybase began to work closely with other vendors, such as HP, with certified configurations, which will mitigate the harm of any loss of interest by IBM.

Teradata

Teradata (www.teradata.com) offers several data warehouse appliances combining hardware, operating system and DBMS. Its offerings include dedicated development boxes, entry-level-priced solutions, data marts and data warehouses, and, since 2010, the first completely solid-state data warehouse appliance (the Extreme Performance Appliance 4600).

Strengths

- Teradata came out fighting in 2010, after other companies' platforms and high-end offerings gained traction in the second half of 2009. It repositioned its midrange data warehouse appliance, the 2650, as a strong competitor to Oracle's Exadata and to Netezza, in a move that gave Teradata two solid offerings. More importantly, this appliance prompts clients to examine both mixed and unmixed workload performance, in which Teradata outperforms its competitors. In late 2010, Teradata announced a partnership with Cloudera that draws on Teradata functionality such as partitioned primary index tables created from a database input format function based on "mappers" that recognize and process MapReduce results.

- Teradata's management software, including Teradata Active System Management (TASM) and Viewpoint, is a clear strength. The management software manages the entire data warehouse environment. Teradata's Analytical Ecosystem Management software is another core strength. It confers the ability to gain a single operational view across Teradata systems and to move and manage data and applications between multiple analytical systems in an enterprise. This software includes Viewpoint, Data Mover, Multi-System Manager, Query Director, Master Data Management and Replication Services. Teradata has a formalized strategy for combining older equipment with new generations ("investment protection"); the use of virtual work units can be distributed, with more work units on newer generation nodes relieving some of the performance pressure on older equipment. In addition to an Enterprise Active Data Warehouse for operational analytics support, features such as object access and query resource filtering, throttles that can be applied to named users, connections or the entire system, and performance groups (high, medium or low priority) contribute to the software's management capabilities.

- Teradata has strengthened its platform's ability to integrate with data mining and advanced analytics by introducing specific support for SAS software. To its support for basic SAS procedures, risk scoring and SAS formats in the DBMS, Teradata has added integration with SAS Enterprise Miner, integrated analytic model management and specific analytics accelerators in the DBMS. Teradata's focus on analytics and data warehousing workloads previously resulted in the introduction of "infrastructure servers" — servers managed within the Teradata cabinet and made available primarily for analytics applications (such as SAS and Viewpoint). With the latest release of the Teradata database, Teradata introduced enhancements in compression and temporal support. The latter provides automated data management and intelligent query processing to enable companies to analyze how conditions have changed or how they looked at any point in time. Teradata also has technical and marketing partnerships with major BI and data integration vendors. Teradata's long experience in the data warehouse space, and its specific strategy for addressing the common use of best-of-breed approaches in the largest organizations, remain **strengths**.

Cautions

- Teradata was invited to participate in more competitive bidding processes in 2010, owing to the growing success of its data mart and data warehouse appliance offerings. This has increased its exposure to competition in the smaller data warehouse markets. Gartner has already noted an unexpected consequence of this, in that some clients report that they selected competitors because there was "no discernible difference" in performance between Teradata's offerings and those of its competitors' appliances in such situations. The same customers also report that their warehouse workloads are either somewhat predictable, or that they have only four or five of the six data warehouse DBMS workloads that Gartner defines. This has forced Teradata to educate prospective customers about the advantages outlined in the **Strengths** section. The net result is that Teradata is competing in more bids but getting mixed reviews from technical evaluation teams when competing with rival vendors' appliance offerings.

- Teradata's customers identify a series of practice and management issues relating to its platform. In general, they focus on the need for Teradata to be more aware of the entire analytics user community that accesses the platform, so that they can derive full value from it. Gartner clients indicate during inquiries that they do not use the platform to its full potential and that its optimization needs are troublesome. The most important issue, however, is that prospective clients are expected to understand the differentiation between Teradata's appliance offerings and the enterprise-class product when deciding on a purchase — and that most entry-level and even second-generation warehouse implementers have difficulty determining the future needs of their users. In short, prospective customers need to be educated about Teradata's approach before they can determine the difference between its products, and, more importantly, between Teradata's appliance and those offered by competitors.

- Common data warehouse practices have renewed the appeal of the single-vendor approach — for example, of buying the complete stack of IBM's Smart Analytics. At the same time, Teradata continues to take a best-of-breed approach. Developments noted above indicate that Teradata is acutely aware of this issue, and its partnerships involve both marketing and technological cooperation. We believe that organizations should focus on decision criteria relating to mixed-workload demands, balanced system management and data optimization, which are pervasive factors in the data warehouse DBMS market and usually more easily addressed by a best-of-breed approach. Nevertheless, single-vendor stack bias has grown for almost three years.

Vertica

Vertica (www.vertica.com) offers a fully integrated column-store analytic DBMS with a number of additional capabilities for high performance and high availability. It derives from research originally done at the Massachusetts Institute of Technology.

Strengths

- Vertica's DBMS has many features that set it apart from other DBMS engines, both column-store and traditional row-store. In addition to using a cluster of commodity servers, giving it scalability and reliability, Vertica's offering has built-in high availability (including active replicas, auto-node recovery and no single point of failure shared-nothing architecture) and data compression (additional to, and different from, the automatic compression realized as a column-store DBMS). Additionally, Vertica has always had a DBMS model using in-memory and on-disk storage in combination, which enables greater performance while maintaining persistence and availability. In 2009, Vertica introduced FlexStore technology, which increases loading and query performance — a benefit substantiated by several reference customers. More recently, Vertica has added the ability to place or "fix" data into flash memory (such as that of Fusion-io), which, in addition to the use of in-memory structures, allows customers to have three levels of storage for hot and cold data. In addition, Vertica was the first DBMS to add connectivity to Cloudera's distribution of Hadoop MapReduce (CDH), enabling users to take advantage of Hadoop MapReduce without implementing it inside the DBMS. It has since added the ability to use the HDFS as a high-speed archive for the data warehouse, and many enhancements to connectivity for better performance when using Hadoop MapReduce.

- Vertica's solution has shown strong adoption as an analytic data mart, gaining more than 280 customers in only a few years, about 20% of which are outside North America. Many Vertica customers have large amounts of SSED loaded into very large databases (a data mart can be huge). Vertica's DBMS is inexpensive, with a pricing model based on the amount of SSED loaded into the DBMS, rather than on the number of users, servers, chips or cores. Its fast adoption is also a result of simple installation and portability across hardware systems. Reference customers report that they can set up Vertica data warehouses very rapidly, sometimes in a matter of hours. This is partly because a feature of Vertica's solution — automatic database design — requires less optimization of the model. This has enabled Vertica's solution to be used by many small organizations with few IT resources.

- Vertica's solution was the first DBMS to run on cloud infrastructure, using Amazon Elastic Compute Cloud (EC2). Thanks to the cloud, implementation and setup are very fast, sometimes taking as little as an hour. Vertica has continued its progress toward stronger cloud-enabled DBMS functionality and true elasticity of cloud service provision (as defined by Gartner) and through partnerships with several cloud infrastructure providers. Not only does this allow for rapid development of analytic applications on Vertica's solution, but it also enables Vertica to perform more POCs and do less travelling to customer sites, while still giving customers full control of their environment through the cloud. In addition to offering stand-alone and cloud versions, Vertica has emphasized multiple implementation models, including SaaS and embedded models (several vendors have Vertica's technology embedded in their software).

Cautions

- There are many competitors in the column-store DBMS space, which makes differentiation difficult. This situation favors mature products with a large installed base, and makes it harder for newcomers like Vertica. Although Vertica does have some differentiation, its challenge is to explain this to prospective customers. Also, like other column-store DBMS vendors, Vertica will face competition from more mature DBMS vendors as they add column-store compression and other capabilities (hybrid column and row store) to their DBMSs. (However, simply implementing a column-store table type in a DBMS is very different from fully integrating all DBMS functionality based on a column-store like Vertica's; this implies that other DBMS vendors will need time to fully integrate and use the column-store model.) Another challenge is posed by Sybase IQ 15.3, which recently added a clustered MPP capability to its DBMS, enabling it to compete better with Vertica.

- Vertica has a few customers with very large data sizes. Because of the exceptional compression in a column-store DBMS, we measure the amount of SSED loaded into the database. Vertica has a few customers with as much as 300TB of SSED, so requiring considerably less storage in the database. Also, according to our reference checks and inquiries, Vertica has only a few customers with large numbers of users (more than 100). Vertica must continue to gain customers with large SSED sizes and greater numbers of concurrent users to compete well against established products, both column-store and traditional. Furthermore, other vendors are adopting Vertica's lower-priced model (based on the amount of SSED loaded into the database) or a hybrid thereof, so reducing one of Vertica's competitive differentiators.

- Vertica has recently added more mixed-workload management capabilities to its DBMS. This enables users to manage better the mix of analytic applications running on the DBMS. However, Vertica's offering is generally not used as a complete data warehouse for multiple subject areas running a mixed workload. For Vertica to compete in the data warehouse DBMS market beyond analytic data marts, it must add more workload management capabilities and the other functions needed for a broader set of data warehouse applications and for managing multiple subject areas. Otherwise, it will be relegated to analytic data mart installations only as other column-store DBMSs increase their capabilities in this area and as mature data warehouse vendors, which already have strong workload management capabilities, add column-store capabilities.

Vendors Added or Dropped

We review and adjust our inclusion criteria for Magic Quadrants and MarketScopes as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant or MarketScope may change over time. A vendor's appearance in a Magic Quadrant or MarketScope one year and not the next does not necessarily mean that we have changed our opinion of that vendor. It may reflect a change in the market and, therefore, changed evaluation criteria, or a change of focus by the vendor.

Acronym Key and Glossary Terms

ASE	Adaptive Server Enterprise (Sybase)
ASM	Automatic Storage Management
BI	business intelligence
CDH	Cloudera's distribution of Hadoop MapReduce
CEP	complex event processing
CPG	consumer packaged goods
DaaS	data warehousing as a service
DBA	database administrator
DBMS	database management system
EC2	Elastic Compute Cloud (Amazon)
ETL	extraction, transformation and loading
FTE	full-time equivalent
HDFS	Hadoop Distributed File System
ICE	Infobright Community Edition
IDE	integrated development environment
IEE	Infobright Enterprise Edition
I/O	input/output
ISV	independent software vendor
MPP	massively parallel processing
OLAP	online analytical processing

OLTP	online transaction processing
PDW	Parallel Data Warehouse (Microsoft)
POC	proof of concept
RAC	Real Application Clusters
RDBMS	relational database management system
SaaS	software as a service
SLA	service-level agreement
SMB	small or midsize business
SQL	Structured Query Language
SSAS	SQL Server Analysis Services (Microsoft)
SSED	source-system-extracted data
SSIS	SQL Server Integration Services (Microsoft)
SSRS	SQL Server Reporting Services (Microsoft)
TASM	Teradata Active System Management
TB	Terabyte
TCO	total cost of ownership

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor that compete in/serve the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability (Business Unit, Financial, Strategy, Organization): Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood of the individual business unit to continue investing in the product, to continue offering the product and to advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all pre-sales activities and the structure that supports them. This includes deal management, pricing and negotiation, pre-sales support and the overall effectiveness of the sales channel.

Market Responsiveness and Track Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message in order to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional, thought leadership, word-of-mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling product that uses the appropriate network of direct and indirect sales, marketing, service and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature set as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including verticals.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.