



Magic Quadrant for Data Warehouse Database Management Systems

23 December 2008

Donald Feinberg, Mark A. Beyer

Gartner RAS Core Research Note G00163473

The data warehouse DBMS market is expanding at a record pace with new vendors, new offerings and high growth. We discuss this, the growth of appliance offerings and how data warehouse DBMS software-only vendors are responding with enhanced functionality and low-cost, market-entry solutions.

What You Need to Know

The data warehouse is a mission-critical system, with data warehouses serving in an increasingly mixed workload capacity, including as a data source for online applications. "Deep mining" analysts and business analysts are running less-structured but equally complex queries and fast-running tactical queries, each with differing service-level expectations. These differing workloads are all competing for CPU, memory and disk access. At the same time, data latency continues to progress from batch to continuous loading demands.

The market has evolved in 2008 in a number of ways. Although this is a mature market with the full attention of large vendors seeking to make inroads with scale and innovation, smaller entrants often deliver a more focused, innovative solution. The latest wave of data warehouse adoption includes less mature organizations with little or no data warehouse management experience. Many of these organizations provide a new opportunity for mart-style deployments. IBM, Oracle and HP have all responded by deploying an appliance solution of some type.

End-user organizations should ignore marketing claims and instead base decisions on customer references and proofs of concept (POCs) to ensure that claims made by vendors will hold true in a real-life environment — more specifically, their own environment.

Organizations are advised to consider their own level of risk aversion when making their data warehouse database management system (DBMS) choice, considering the following:

- With most of the market's major vendors providing both software-only and appliance-based solutions, it is possible to select a solution offered by almost any major vendor that is also an organization's preferred or standard vendor. However, specific advantages in managing larger data volumes and mixed workloads should be

Acronym Key and Glossary Terms

| | |
|--------------|--|
| ASM | Automatic Storage Management |
| BI | business intelligence |
| BW | Business Information Warehouse |
| DBA | database administrator |
| DBMS | database management system |
| EDW | enterprise data warehouse |
| ETL | extraction, transformation and loading |
| I/O | input/output |
| MPP | massively parallel processing |
| OLAP | online analytical processing |
| OLTP | online transaction processing |
| POC | proof of concept |
| RAC | Real Application Clusters |
| RDBMS | relational database management system |
| SaaS | software as a service |

considered, as should the increasing importance of managed support. Most of the leaders represent low-risk options, even if the selected platform is outside enterprise standards.

- Niche vendors are offering new technology in the market, such as data tokenization, the use of specialized compression techniques and data storage techniques that reduce input/output (I/O) constraint issues for high performance. These vendors include risk aspects such as small reference bases (and commensurate lack of best practices) and vendor capitalization issues. Organizations that embrace a higher level of risk should seek compensation in the form of discounts, and even extended pilot implementations with below-average prices for support during the pilot period.
- Some of the largest software and hardware vendors now play a new role offering new and unproven products. The risk to organizations here is in these vendors' commitment to future product development. Before committing, organizations should evaluate the vendor road map and review its track record of meeting product development objectives.
- In one of the most interesting advances in 2008, the leading and challenging vendors advanced various solutions to meet functionality and packaging demands in the market. At the same time, the market matured enough to realize the benefits of existing functionality in some of the platforms. The market demanded low-entry-price solutions that scale up and out as part of a cost-control strategy. The demand for data mart support increased and vendors responded with summary, dimensional and multidimensional deployment support. Mixed workloads began to overwhelm enterprise warehouse performance and solution providers responded with synchronized, dual-warehouse deployments. These and other vendor responses had the effect of reducing the "white" space at the top right of the Leaders' quadrant as the vendors "covered" the demands with solutions.

[Return to Top](#)

Magic Quadrant

Figure 1. Magic Quadrant for Data Warehouse Database Management Systems

| | |
|-------------|---------------------------------|
| SQL | Structured Query Language |
| SSAS | SQL Server Analysis Services |
| SSIS | SQL Server Integration Services |
| TCO | total cost of ownership |


Note 1 Definition of Mission-Critical Systems

Mission-critical systems are defined as systems supporting the generation of revenue or supporting business processes 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 increased business costs. Normally, mission-critical systems require high-availability systems and disaster recovery sites. We have included 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 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 forces issues in summary and aggregate management to support dashboards and prebuilt reports.
- Batch data loading continues to persist as the market matures and begins to realize that not all data is required for "right time" latency and that some information, being less volatile, does not need records refreshed as frequently as the more dynamic real-time data elements.
- Large numbers of standard reports ranging in the thousands per day requiring SQL tuning, index creation, new types of storage partitioning and other types of optimization structures 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 (pre-joins) and designated dimensional drill downs (summary). They rely on a BI architect to develop commonly used cubes or tables.
- An increasing number of true ad hoc query users (data miners) with a random, unpredictable use of the data, implying a lack of ability to specifically tune 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 the OLTP applications requiring high-



Source: Gartner (December 2008)

[Return to Top](#)

Market Overview

The data warehouse DBMS market has evolved from a traditional information store supporting business intelligence (BI) users and tools into an analytics infrastructure repository for the enterprise. And organizations are adding additional workload from online transaction processing (OLTP) applications and increasing the frequency of data loading to intraday, approaching continuous loading. The data warehouse DBMS market expanded in 2008 to include new areas seeking to extend existing data marts into broader data warehouses or seeking some level of federation for existing low-cost solutions. This has created opportunities for a rebirth of appropriately less sophisticated market-entry solutions, but also the opportunity for a resurgence in discredited, archaic practices for the inexperienced end-user organization. At the high end, data warehousing is now mission-critical (see Note 1). In addition, the size of the database is becoming less important. In the past, buyers believed that the vendor with the largest database was the leader. Today, smaller data warehouses (those less than 5TB) are commonly solving organizations' analytic needs. Other criteria — notably cost, given the current economic conditions — are more important than database sizes.

We continue to see new vendors entering the market and mature vendors offering new solutions. In 2009, we will be watching several new vendors of DBMS engines that currently meet some, but not yet all, of the criteria required to be included here — Aster Systems, EnterpriseDB's GridSQL and ParAccel. The lessons learned from DATAAllegro, Netezza and Teradata about the advantages of preconfigured, balanced appliance solutions (see "Data Warehouse Appliances Are More Than Just Plug-And-Play") caught the interest of some more traditional vendors, such as HP (Neoview), IBM (InfoSphere Warehouse), Oracle (HP Oracle Database Machine) and Sun Microsystems (using the Greenplum DBMS with the Sun Fire X4540 data server).

performance queries. This is one force driving the requirement of high availability in the data warehouse.

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 appearing in a Magic Quadrant or MarketScope one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. This may be a reflection of a change in the market and, therefore, changed evaluation criteria, or a change of focus by a vendor.

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 and skills, 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 presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support and the overall effectiveness of the sales channel.

[✦ Return to Top](#)

Data Warehouse Mixed Workloads

The traditional data warehouse workload of queries and reporting is changing to a data warehouse with a mix of several distinct workloads (see Note 2). These six workload types are creating more issues for vendors than the actual size of the data warehouse, even manifesting in databases smaller than 1TB. In addition to service-level expectations (see "Mission-Critical Data Warehouses Demand New SLAs"), the size and duration of "useful" data for each community often differs significantly, forcing every aspect of the data warehouse environment to become involved — from I/O channel balancing, through disk management, and into memory and processor allocation. Through 2011, mixed workload performance will remain the single most important performance issue in data warehousing. As a direct effect of the complex mixed workload, with continuous loading and the increase in automated transactions from the functional analytics in OLTP, the transactional DBMSs may be able to erode the performance edge that was formerly attributed to specialized data warehouse DBMS solutions.

[✦ Return to Top](#)

The Intensive POC

Performing a POC with a "shortlist" of vendors during the selection phase of the data warehouse infrastructure has become a best practice. This is especially important when considering one or more of the newer entrants to this market. We recommend POCs using as much real source system extracted data from the operational systems as is reasonable to use. We also recommend performing the POC with as many users as possible, creating a data warehouse workload approaching the environment to be used in production. Do not provide some of the more complex queries to vendors in advance of the POC to be certain the DBMS has not been "pretuned" for your queries. We recommend data loading as part of the POC, even if that is not one of the important requirements of the system. If continuous loading is a requirement, or batch loading is desired, they must be part of the POC. Understanding a solution's capabilities and/or restrictions is important, as the size of the window for loading in most organizations is diminishing as the data warehouse becomes a 24/7 operation.

[✦ Return to Top](#)

Optimization and Performance

Sophisticated data warehouse platforms are now the rule, rather than the exception. Such platforms include hardware management of I/O, disk storage and CPU/memory balancing almost as a matter of course. However, new customers to the market are often ignorant of the issues that differentiate analytic processing optimization from transactional optimization. The same education curve that has been the mainstay of the data warehousing industry is now taking place in microcosm in these new entrant organizations. Data warehouse novices are advised that this is a mature industry, which means there are highly qualified and capable solution providers, architects, database administrators (DBAs) and modelers. It also means there is a significant opportunity for disingenuous implementers to sacrifice appropriate design for the age-old time-to-market promise. One of the favorite tactics of less capable implementers is to defer optimization considerations until issues are encountered. Organizations are cautioned to consider that all data warehouse implementations will encounter optimization requirements — usually within 18 months of their initial deployment. The time frame for this issue is such that it negates most time-to-market justification reasons.

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 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 and service-level agreements.

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 Web site, 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.

IBM, Netezza, Oracle, Sybase and Teradata have all introduced various new styles of optimization in the past 15 months. Their long experience in data warehousing has placed them in an ideal position to identify "grassroots" issues and deploy solutions designed to meet real-world situations (for example, IBM's compression, Oracle's incremental updates to materialized views and compression, Netezza's zone maps and Teradata's hot and cold data placement). The new entrants are focusing on optimization as a differentiator; for example, column stores (as opposed to traditional row stores), tokenization of data, and hardware parallelization. Finally, nearly every data warehouse vendor is now addressing the issue of optimized storage for the warehouse, not only using the concepts of hot and cold data, but also by using different sizes of storage devices to reduce cost (larger drives with lower performance) or increase performance (smaller drives at a higher cost per terabyte). In addition, many vendors are moving DBMS code into, or closer to, the storage devices, gaining a higher degree of parallelism and more computer power (by using the processors in the storage devices).

[✦ Return to Top](#)

A Wave of Data Warehousing Novices

In 2008, an adoption wave by data warehousing novices — as well as more conservative organizations with highly skilled personnel — entered the demand side of the market. The result of this wave of adoption initially placed opportunistic pricing pressure on the leading vendor solutions. However, in the past 15 months, nearly every vendor has introduced new, entry-point solutions that have adequate performance for lower cost (for example, Teradata's low-cost appliances, Microsoft's simple licensing model, IBM's Linux InfoSphere Balanced Warehouse and Netezza's capacity-on-demand licensing). The vendors were also careful to provide a scale up/out strategy to their customers if these entry solutions were used. The result is that the traditional market-leading vendors seem to have started their perimeter defense by introducing lower price points, creating new entry-level solutions and adding additional features without additional pricing.

Another effect of this wave of data warehouse demand and adoption is a resurgence in some older, but still effective, data warehousing practices. While these aging practices have not returned to their former status of "best" practice, they certainly represent an opportunity for vendor marketing. The deployment of data marts or data mining platforms is acceptable again. The use of stand-alone data marts remains a common practice in the marketplace, but most organizations now realize that if these data marts are properly positioned as a basic solution, they can eventually be retrofitted or rearchitected as part of a more comprehensive strategy. In 2008, this marketing vision translated into execution and, more specifically, enhanced revenue. When an organization is initiating a data warehousing strategy and has nothing, having a well-architected data mart with plans to expand it into a larger enterprise data warehouse is a valid strategy. It is important to note the difference between this strategy and the previously failed strategy of independent data marts deployed as a replacement for a comprehensive data warehouse strategy. Additionally, dedicated analytic workload platforms are also making a resurgence, as more mature data warehousing teams realize the need to isolate the more intensive workloads, such as data mining.

[✦ Return to Top](#)

Alternative Delivery Models

- *Managed service warehouse.* The use of data warehouses as a managed service has been an option in this market for more than 10 years but, in 2008, at least two vendors focused on managed data warehouses as a business model (1010data and Kognitio). The DBMS

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.

vendor would develop and run the data warehouse for the customer, with the customer's business unit purchasing a managed service instead of using its own data warehouse. Normally our references reported that the IT organization did not have the bandwidth to support the specific application because of overloading of the data warehouse and/or the IT staff — thus the IT staff supported the effort. Now, many vendors offer some form of managed data warehouse. Pursuing managed data warehouses exclusively is no longer a viable business model, but it remains a valuable alternative channel. We expect the use of this model to increase in the next few years, especially for single business units or a specific data warehouse application. Further, we believe this model will develop into a software as a service (SaaS) model in the next few years for organizations in the small or midsize business category that lack the expertise and funds to support their own data warehouse.

- *Open-source data warehouses.* Open-source DBMSs are now being used in experimental and more formalized approaches. At this point, open-source warehouses are rare and usually smaller. They also generally require a more manual level of support. However, some solutions, such as Greenplum — commercial software using an open-source DBMS (PostgreSQL) — are optimized specifically for data warehousing.
- *Resurgence of data marts.* In 2008, organizations showed an increase in mart-style implementations. The prompt response by some vendors (notably Sybase) to capitalize on this shift quickly translated from a marketing vision into execution in 2008. The resurgence of the data mart has pros and cons. Because of the wider acceptance and increasing variety of data-warehouse-driven applications and workload variations, data mart proliferation has become a primary concern. A data mart is defined as an application-specific analytic repository of any size, normally with a specific, smaller group of users than an enterprise data warehouse (EDW) (see "Of Data Warehouses, Operational Data Stores, Data Marts and Data 'Outhouses'"). An application can be a traditional set of applications (for example, SAP) or a workload-specific application (such as business analysts using a subset of the data for data mining purposes). The use of data marts specifically for analytics continues with unabated growth. This is not only because of the workload that analytics can place on the EDW, but also because some of the data warehouse DBMS engines excel in analytic application. Data marts can be used to optimize the EDW by offloading part of the workload to the data mart, returning greater performance to the warehousing environment. Specifically, column-oriented DBMS engines — such as ParAccel, Sand/DNA Analytics, Sybase IQ Analytics Server and Vertica — have shown excellent performance in analytic applications compared with the more traditional row-based DBMSs. Several references reported an improvement in performance of up to 100 times. However, when used in a complex query, with many columns with complex joins, column-based DBMSs sometimes perform no better, and on occasion worse, than a row-based DBMS. If considering a column-store DBMS as an EDW, a POC remains mandatory.
- *Distributed data warehouse.* These warehouses continue to see adoption, but at a very measured pace (see "Emerging Trends: Introducing the Distributed Data Warehouse"). Reasons given for this approach vary from the creation of physical data security zones to global operations requiring 24/7, time-zone-based analytics. The approach should not be confused with a federated design.

✦ [Return to Top](#)

Market Definition/Description

The data warehouse DBMS market consists of the vendors supplying DBMS products that provide the database infrastructure of the data warehouse.

For the purpose 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 those systems that, in addition to supporting the relational data model (extended to support new structures and data types such as materialized views and XML), also 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. It is important to note that a DBMS cannot "be" a data warehouse. It is the platform on which a data warehouse (solution/data architecture) is deployed. This market is specific to DBMSs that are used as a platform for a data warehouse.

A data warehouse is one 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, though Gartner defines a *small* data warehouse as less than 5TB, a *medium* data warehouse as 5TB to 50TB, and a *large* data warehouse as greater than 50TB. For the purposes of measuring the size of a data warehouse database, we define data as *source system extracted data*, excluding all data warehouse design-specific structures (such as indexes, cubes, stars and summary tables). Source system extracted data is the actual row/byte count of data extracted from all sources.

[✦ Return to Top](#)

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 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 consider open-source DBMS products from vendors that control or participate in the engineering of the DBMS (see "The Growing Maturity of Open-Source Database Management Systems").
- Data warehouse DBMSs or DBMS products that support an integrated front-end tool, but which can also open their DBMS to competing applications, are included if access is achieved via open-access technology, as opposed to custom-built application programming interfaces (APIs).
- 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 the front-end tools, or that exclusively support an integrated front-end tool, do not qualify for this Magic Quadrant.
- The Magic Quadrant is based on vendors. Where a vendor supplies more than one distinct DBMS for data warehousing, our evaluation considers all products together for analysis and specifies where there are differences.

[✦ Return to Top](#)

Added

- 1010data.
- HP.
- illuminate Solutions.
- Ingres.
- Vertica.

[✦ Return to Top](#)

Dropped

- DATAlegro was acquired by Microsoft in August 2008 and is now merged with the evaluation of Microsoft.

[✦ Return to Top](#)

Evaluation Criteria

Ability to Execute

Ability to execute is primarily concerned with the ability and maturity of the product and the organization. These criteria also consider the portability of the product and its ability to run and scale in different operating environments, giving the customer a range of options. This also includes the differentiation between data warehouse DBMS solutions and data warehouse appliances. The ability to execute criteria are critical to the level of satisfaction and success the customer has attained with the product, and so customer references are weighted heavily throughout these criteria.

Specific Criteria:

- **Product and service** includes the technical attributes of the DBMS. We include scalability, manageability, security, high availability/disaster recovery, support of mixed workloads, support of additional data structures (such as XML) and data loading. These attributes are measured across a variety of database sizes and workloads. Also, we consider the resources necessary to manage the data warehouse, especially as the data warehouse scales to larger sizes and more complex workloads.
- **Overall viability** includes the corporate aspects of ability to execute, such as the skill level of the personnel, financial stability, R&D investment, and merger and acquisition activity. This also includes management's ability to be responsive to market changes and, therefore, the ability of the company to survive through market difficulties (critical to the long-term survival of the vendor).
- Under **sales execution and pricing**, we examine the price and different pricing models of the DBMS, the ability of the sales force to manage accounts and whether the sales team is compensated appropriately in line with the corporate marketing initiatives. We also include the channel partnerships here, and the ability of the vendor to create and use the partner model.
- **Market responsiveness and track record** covers the issue of references (for example, how many, what size of companies, what configurations and workload mix). Also included is the ability of the vendor to adapt to market changes and its history of being flexible to market dynamics.
- **Marketing execution** explores how well the vendor understands and builds its products in response to customers' needs, in addition to targeting offerings to these needs and to the needs of the market in general. This criterion includes the completeness of the vendor's offering as well.
- Customer support and professional services are evaluated as part of the **customer experience** criterion, together with input from customer references as described earlier. Also included is the track record of POCs and customer perceptions of the product, as well as

aspects of customer loyalty to a given vendor. This demonstrates customer tolerance of vendor practices and may indicate satisfaction.

- **Operations** covers the alignment of the company's operations, as well as whether and how they enhance the ability of the company to deliver.

Table 1. Ability to Execute Evaluation Criteria

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

Source: Gartner

[✦ Return to Top](#)

Completeness of Vision

Completeness of vision encompasses the ability of the vendor to understand the functionality necessary to support the data warehouse workload design, the product strategy designed to meet market requirements, and the ability to understand overall market trends and influence or lead the market when necessary. A visionary leadership role is necessary for the long-term viability of the product and the 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. However, Gartner clients are cautioned to be wary of vendors with extremely good vision (including the communication of that vision) but with low execution capability. Data warehouses are mission-critical — and poor execution will begin to hurt the overall viability of the organization.

Specific Criteria:

- **Market understanding** covers the ability of the vendor to understand and shape the data warehouse DBMS market and show leadership in it. In addition to examining the core competencies of the vendor in this market, we also consider the awareness of the vendor of new trends in the market.
- **Marketing strategy** refers to the vendor's marketing messages and its ability to choose appropriate target markets and third-party software vendor partnerships to enhance the marketability of its products. For example, does the vendor encourage and support independent software vendors (ISVs) in its effort to support the DBMS in native mode?
- An important criterion for vision is the **sales strategy**. This encompasses all the channels and partnerships developed to assist with sales. This is especially important for younger organizations, allowing them to greatly increase their presence in the market while maintaining a *lower cost of sales*. This criterion also includes the company's ability to communicate its vision to its field organization and, therefore, to clients and prospects.
- **Offering (product) strategy** covers the areas of portability and

packaging of the products. Vendors must demonstrate a strategy that enables customers to choose what they need to build a complete data warehouse solution. We also consider the availability of the vendor's DBMS as a data warehouse appliance.

- The **business model** covers how the vendor's model of a target market combines with product offerings and pricing, and whether it has the ability to produce profits with this model based on the 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 the ability of the vendor to understand its clients. Specific models for the data warehouse, however, belong in a discussion of applications.
- **Innovation** is a major criterion for evaluating the vision of data warehouse DBMS vendors in developing new functionality, R&D spending, pushing the market in new directions and "pushing the envelope" in the market. This also includes the vendor's ability to innovate and develop new functionality in the DBMS, specifically for the data warehouse. Increasingly, users are expecting the DBMS to become more self-managing and self-tuning, reducing the resources involved in optimizing the data warehouse, especially as the mixed workload increases.
- The organization's worldwide reach and **geographic strategy** are evaluated considering the organization's ability to leverage the resources in geographic regions, as well as subsidiaries and partners in other regions. This is becoming increasingly important, as the number of regionally distributed data warehouses increases (as discussed in the Market Overview section). A vendor's success increasingly depends on its ability to market and support its data warehouse DBMS in a geographically dispersed area, using subsidiaries or distributors. This criterion also includes the ability of the vendor to support clients throughout the world, around the clock, in many languages.

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 | standard |
| Vertical/Industry Strategy | standard |
| Innovation | high |
| Geographic Strategy | standard |

Source: Gartner

[✦ Return to Top](#)

Leaders

The Leaders' quadrant for data warehouse DBMSs contains those vendors that demonstrate the greatest degree of support for data warehouses of all sizes, with large numbers of concurrent users and management of mixed data warehousing workloads. These vendors lead the market 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. Because of this track record, leaders also

represent the lowest risk for successful data warehouse implementations. Additionally, the maturity of this market demands that leaders maintain a strong vision regarding the key points emerging during the past year: mixed workload management for end-user service-level satisfaction and data volume management.

[✦ Return to Top](#)

Challengers

In the past, the Challengers' quadrant has represented vendors with strong offerings for the client base. In 2008, Gartner clients report that vendor stability coupled with an established offering constitute a challenger. These vendors have market presence in the data warehouse DBMS space, a proven product and have demonstrated corporate stability. Challengers generally have a highly capable execution model. Ease of implementation, clarity of message and end-client engagement all contribute to making these vendors successful. Challengers show a wide variety of data warehousing implementations across different sizes of data warehouses with mixed workloads. Organizations often purchase the products of challengers for deployment in a limited fashion, such as a departmental warehouse or large data mart, with the intention of scaling the solution to enterprise class. The products usually encounter resistance as an enterprise solution because they vary from corporate standards or the market is uncertain of the vendor's commitment to the space. These vendors exhibit the potential to move into the Leaders' quadrant by demonstrating strong, new client acceptance as an enterprise solution.

[✦ Return to Top](#)

Visionaries

Visionaries take a forward-thinking approach to managing the hardware, software and end-user aspects of the data warehouse. Visionaries frequently suffer from a lack of global, or even strong regional, presence. They normally exhibit a smaller market share than leaders and challengers. New entrants with exceptional technology may appear in this quadrant very early after their products have become generally available but, more typically, vendors with unique or exceptional technology will appear in this quadrant when their products have been generally available for several quarters. The Visionaries' quadrant is often populated by new entrants that have new architectures and functionality that are unproven in the market. Vendors must demonstrate customers in production proving the value of the new functionality and architecture. The requirement for the existence of production customers and general availability of at least a year indicate that visionaries must be more than a startup with a good idea. Frequently, visionaries will drive the leaders toward new concepts and engineering enhancements.

[✦ Return to Top](#)

Niche Players

A niche player has low market share or low market appeal. Frequently, a niche player provides an exceptional data warehouse DBMS product, but it is isolated or limited to a specific end-user community, a specific region or a specific industry. Although the solution itself may be without limitations, market adoption is limited. This quadrant contains vendors in two categories: smaller vendors with data warehouse DBMS products that lack a strong or large customer base, or smaller vendors with a data warehouse DBMS that lacks the functionality of those of the leaders. Niche players typically offer smaller, specialized solutions that are used for specific data warehouse applications depending on the needs of the client. This quadrant

also includes vendors of new data warehouse DBMS products that lack general customer acceptance or the proven functionality to move beyond niche status. This is the starting point for many new entrants.

[✦ Return to Top](#)

Vendor Strengths and Cautions

1010data

This is a seven-year-old managed service data warehouse provider, based in New York City, with an integrated DBMS and BI solution targeted at the business side of organizations, primarily in the financial sector. As such, it has a different operating model from the other vendors in the Magic Quadrant.

[✦ Return to Top](#)

Strengths

- 1010data offers a solution including a DBMS to provide high-speed analytics for businesses. As 1010data is a SaaS vendor, the business unit needs little experience in data warehousing or BI. Also, as 1010data is a managed service solution, it can complement the internal IT department with fast-to-market solutions for business units, alleviating the resource consumption within IT.
- A big challenge for data warehouse SaaS solutions is remote locations and issues with security and data transfer performance (perceived or real). 1010data can host its solution for its customers in the traditional SaaS model. However, the vendor also offers a managed solution at the customer site, removing these issues.
- The managed service model allows 1010data to leverage software solutions across multiple customers. This increases the availability of applications to businesses. The best example of this is connected to the recent banking issue concerning subprime loans. 1010data was able to implement a risk analysis model quickly for many institutions.
- Its product is a good, fast-to-market solution for organizations needing a short-term BI application, or lacking BI and data warehousing expertise. The solution is a DBMS that is fully compliant with Structured Query Language (SQL) and has an Open Database Connectivity (ODBC) interface that can be used for other applications, in addition to its own.

[✦ Return to Top](#)

Cautions

- With only a fully managed service model, 1010data is susceptible to push-back from IT departments that wish to have all data warehouses in-house with governance of the data asset of the organization.
- 1010data is an integrated DBMS and BI solution and, therefore, not a pure DBMS choice.
- As a small vendor with a small sales force, 1010data is more susceptible than bigger vendors to DBMS market fluctuations, which means there is a higher risk in using its solution.

[✦ Return to Top](#)

Greenplum

Greenplum is a massively parallel processing (MPP) data warehouse DBMS based on open-source DBMS PostgreSQL running on Linux and Unix. It is sold primarily through channels as an appliance offering (by, for example, Sun Microsystems, which accounts for half of Greenplum's customers).

[✦ Return to Top](#)

Strengths

- Greenplum is a stand-alone DBMS engineered for data warehousing. It has demonstrated scalability in production to hundreds of terabytes and internally to more than a petabyte (1,000TB). It has also demonstrated the ability to run and manage the mixed workload in a number of references. Through its software architecture, Greenplum is able to move DBMS code to the storage device, thereby increasing performance. Greenplum is one of the first data warehouse DBMSs to implement MapReduce for high-scale analytics.
- Since 2006, primarily through partners and channels, Greenplum has acquired more than 50 customers globally. Initial customers report performance at least equal to that of Teradata's solution in a live application environment, at a fraction of the cost.
- A strength of Greenplum's is its strategy of selling through partners and system integrators (such as Sun Microsystems, which sells an appliance with the Greenplum DBMS and the Sun Fire X4540). This saves it from having to build a sales force — an expensive proposition for a small startup vendor — and allows it to concentrate on development and support. As in the case of Sun Microsystems, the partners provide the balanced, packaged configuration of hardware with the Greenplum DBMS as an appliance offering, along with the single point of service required of an appliance.
- The company's use of an open-source DBMS as the core work engine also helps to reduce costs, while it concentrates on the management software surrounding the data warehouse and the optimization features necessary for a complex, mixed workload environment.

[✦ Return to Top](#)

Cautions

- Although some customers have implemented non-Sun systems, Greenplum must add additional major partners and not depend solely on Sun Microsystems. If it does this and then loses a partner, the impact on revenue will be minimized. Further, we expect to see the major hardware vendors (including Sun) having more than one data warehouse appliance. The loss of a major partner would put a strain on Greenplum's resources as a small company.
- As competition in data warehouse DBMSs grows and matures, especially from the traditional vendors such as IBM, Oracle and Teradata, Greenplum will have increasing difficulty differentiating itself in the market. In several years, Greenplum may be acquired by a vendor in the market to complement its technology, much as Microsoft acquired DATAlegro in 2008.

[✦ Return to Top](#)

HP

HP introduced Neoview to the market in early 2007 after carefully cultivating reference customers to assist in the product launch. HP Neoview debuts in the Challengers' quadrant and needs to demonstrate market acceptance in 2009 to remain high in this analysis.

[✦ Return to Top](#)

Strengths

- HP Neoview has near-linear scalability in processing capacity and data storage management, and judging by the comments of references, the overall system architecture can address data warehouse sizes

greater than 20TB.

- HP Neoview optimizes all available memory when executing large queries. This is very effective when very large tables need to be joined to many small tables, which is the usual configuration for a data mart or star schema. As a result, HP Neoview can reduce the amount of optimization work required in larger data warehouses. The functionality is similar to "broadcasting" smaller tables as copies — but it does it in memory to avoid actually making the copies on disk, which would increase storage requirements.
- HP's acquisition of Knightsbridge Solutions in 2006 seems like old news now. However, the field experience provided by this consulting organization offers significant data warehouse implementation experience and enhances HP's ability to deliver implementation services, including data architecture and sourcing strategies. HP gained a successful methodology and experience base from its acquisition of Knightsbridge. As this experience base includes many of HP's current competitors, this allows it to evaluate issues arising from the retrofit of existing warehouses or a "rip and replace" strategy.
- Although Neoview is young and lacks the maturity of the solutions of incumbent providers, HP's strength in the breadth of its resources will help to reduce the risk associated with a relatively new product in the current economic crisis.

[✦ Return to Top](#)

Cautions

- While HP's product is in place with multiple reference customers, many require further confirmation of the solution's maturity before placing it in a trusted production role. Neoview does have references that indicate the system is in production, but further examination usually reveals multiple qualifying statements regarding the definition of production. HP references often report that they have not placed HP Neoview into production and it remains in pilot status for long periods.
- HP has embarked on a mission that addresses greenfield as well as competitive replacement with equal vigor. Neoview is being positioned as a high-end data warehouse platform targeted at CIOs and technical stakeholders with a technology-oriented value proposition. The largest part of the current warehouse market needs a more direct message that does not require advanced warehousing experience to determine the value of the Neoview solution. The apparent complexity of the messaging, combined with a low number of reference customers, poses too many questions, and prospects become hesitant in favor of more tenured solutions in the market.

[✦ Return to Top](#)

IBM

IBM remains a leader in data warehousing DBMSs. In this highly competitive, mature market, IBM has positioned its offering as a full solution in data management where the warehouse is a key part of the strategy.

[✦ Return to Top](#)

Strengths

- IBM can engage organizations that desire a preloaded solution or those that want to build their own hardware environment. IBM's InfoSphere Warehouse (a data warehouse offering based on IBM DB2) is a software-only solution. Its data warehouse appliance solution, the IBM InfoSphere Balanced Warehouse, is a combined

server and storage hardware solution (using the System p server with AIX, or the System x server with Linux or Windows and the IBM InfoSphere Warehouse) complete with service and support.

- IBM InfoSphere Warehouse includes many data warehousing features, such as embedded analytics, data visualization and transformation capabilities, integration with SAS and SPSS, logical and physical data partitioning, compression, workload and performance management, and multitemperature warehouse support.
- IBM has hundreds of customers running medium data warehouses (up to 50TB of source system extracted data), primarily the System z and System p. IBM also has customer references running data warehouses with IBM InfoSphere Warehouse for large amounts of source system extracted data (see Market Overview section). This large customer base also provides significant experience for IBM's professional services offering.

[✦ Return to Top](#)

Cautions

- IBM offers a complete data management solution including middleware, data architecture tools and DBMS-based warehouse solutions. As a result, the IBM Software Group continues to face the issue of "co-opetition" (cooperative competition). The hardware divisions not only sell hardware for other DBMS platforms, but are also beginning to openly support other DBMSs, and even to partner with other DBMSs with competing appliance offerings. This has the effect of IBM confusing its customers and, in some instances, driving customers to examine alternatives to IBM products, lengthening sales cycles and, sometimes, losing customers.
- IBM has introduced new product names (for example, IBM InfoSphere Warehouse, formerly known as DB2 Warehouse Edition) and new configurations and offerings (for example, IBM InfoSphere Balanced Warehouse). While IBM manages to reconcile potential contractual issues, this does not alleviate a customer of its own responsibility to validate that any potential issues have been reconciled for internal control. End-user organizations should also review plans for new IBM products as they relate to organizational two-, five- and 10-year plans.
- In the data warehouse DBMS market, IBM sells almost exclusively to its software infrastructure installed base, rather than capturing new customers. The exception to this is primarily driven by new customers coming from the purchase of SAP applications hosted on DB2. IBM must focus on new data warehouse DBMS implementations increasing its customer base or it will continue to grow at rates lower than those of the overall market (see "Market Share: Relational Database Management System Software by Operating System, Worldwide, 2007").

[✦ Return to Top](#)

illuminate Solutions

A small software vendor based in Spain, illuminate Solutions has an integrated data warehouse DBMS (correlation) and BI tools.

[✦ Return to Top](#)

Strengths

- The company's solution is a unique system that stores data and indexing in what is termed a "data model independent database." The focus of the system is to store all the potential relationships between any data element in the database with any other data element. This

structure is automatically built and maintained via illuminate's DBMS intellectual property as data is loaded — DBAs that are used to row- or column-vectorized, hierarchical files, or any other data file management system, do not have to develop custom load processing.

- Query processing is enhanced, as this solution effectively creates prejoins for all existing data relationships in the data model. The process is repeated when new data sets are added with an almost "spider web" effect that stores every correlation that can be inherited from the data already stored in the database, as well as from any of the newly added data.
- The solution stores abstracted data values as a metadata master set in the database. This enhances data quality by ensuring single storage of each value. A purely column-vectorized approach reduces the volume of the database, as repeated values within a column are addressed, but repeated values in the overall database are still possible. The correlation theory and its use of metadata eliminate even those remaining multiuse redundancies.
- Customers report that illuminate's solution is easy to use and has very high performance. This is to be expected when performing queries of an analytic nature.
- The company has more than 50 customers, primarily in Spain and Europe, with a few in the U.S. and Latin America.

[✦ Return to Top](#)

Cautions

- Data volume is not the only performance inhibitor for data warehouses. The illuminate solution is highly dependent on processor and memory capacity. Currently, Moore's law is a leverage point for this type of solution because, from a speed-of-processing perspective, CPU and memory far outpace disk and I/O speeds. Even in the future, this approach will continue to have advantages. However, the engineering involved will serve to change the equation in illuminate's case. The more efficient data storage and reduced size of the I/O requirement that results will enable the storage system to overwhelm subpar CPU/memory plants.
- The vendor has little presence outside Europe, with most of its business in Spain. Although it has opened an office in the U.S., there has been little activity in that market so far.
- In 2008, the company created its partner program for channels and third-party software. This is an issue that illuminate needs to address in 2009 and beyond.

[✦ Return to Top](#)

Ingres

Ingres' solution is now an open-source, general-purpose DBMS with a 30-year history as one of the original relational database management system (RDBMS) engines. The company has many customers running mission-critical applications, including data warehouses.

[✦ Return to Top](#)

Strengths

- Ingres, a mature vendor in this market, has more than 10,000 customers using its DBMS. Most customers have OLTP applications, but Ingres has its share of smaller data warehouses (less than 1TB to 2TB). Ingres has converted virtually all its pre-open-source customer base to open-source subscription support.
- Ingres has been gaining some third-party software partners, specifically in the BI market. Open-source vendor JasperSoft has a

software appliance or bundle with Ingres for BI. This is driving a larger number of installations in data warehousing, with both new and existing customers looking for an open-source stack supporting BI.

- The company's solution is the only open-source DBMS with proven maturity in mission-critical applications, including data warehousing.

[✦ Return to Top](#)

Cautions

- Data warehousing is not a strength for Ingres, with its solution used mostly in OLTP applications. Ingres' data warehousing professional services are limited. Ingres must also address the areas of enhanced data warehouse functionality, storage management and mixed workload management to compete in the data warehousing DBMS market.
- The history of Ingres works against it, as it is 30 years old and has not regained major market traction. This is an issue with market perception, which is difficult to change.
- While owned by CA, and before becoming a separate open-source company, Ingres gained few new customers, contributing to its negative market perception. Although Ingres has gained new customers and new third-party relationships since becoming an open-source company, it must continue to show increased growth in the number of new customers and in revenue to become a serious competitor in this market.

[✦ Return to Top](#)

Kognitio

Several years ago, Kognitio began by offering data warehousing as a hosted service. Today, it has a mix of customers using its DBMS separately as a data warehouse DBMS engine, as well as using data warehousing as a managed service, hosted on hardware located at Kognitio's site.

[✦ Return to Top](#)

Strengths

- Kognitio comes from a strong DBMS appliance background (following its merger with WhiteCross Systems) and has a track record of solid performance. References report that performance is excellent, with large database sizes in an analytics and reporting environment for thousands of users.
- Kognitio's solution is the original data warehouse DBMS to be primarily used as a managed service provided by the DBMS vendor, with its clients buying their data warehousing services from Kognitio, while Kognitio hosts the database. Recently, we have seen more activity in the market with this model. There are two reasons for this. First, some business units are dissatisfied with their IT department. Buying a managed service for their BI needs bypasses the IT department. Second, some small companies cannot afford their own data warehouse infrastructure but have large volumes of data to process for analytics. This is a growing market.
- Kognitio is selling this model to other vendors wanting to offer the same type of managed services. We believe that, in 2009, we will see additional sources of managed services on Kognitio's solution (with the added advantage of Kognitio breaking out of the Europe-only market).
- In 2007 and 2008, Kognitio gained several large clients that installed its system on-site, as opposed to taking it as a managed service. These installations are large, analytic data warehouses. This demonstrates Kognitio's ability to supply a data warehouse DBMS

capable of competing with those of many of the market incumbents. In addition, Kognitio offers the customer the added flexibility of moving to or from a hosted model as desired.

[✦ Return to Top](#)

Cautions

- Kognitio remains primarily a Europe-only vendor, with most of its customers in the U.K. It has opened offices in the U.S. and is beginning to acquire customers in North America.
- Although Kognitio offers its data warehouse DBMS for sale, it is usually considered for a managed service solution. As a stand-alone solution, it must compete with the incumbents (such as IBM, Netezza, Oracle and Teradata), which is becoming increasingly difficult in this market.

[✦ Return to Top](#)

Microsoft-DATAIlegro

Microsoft is surging in the market because of its low entry price and changes in market demand, which is satisfied with adequate capability rather than advanced functionality. Microsoft's acquisition of DATAIlegro promises a DBMS for data warehousing in the future that will leverage hardware scalability through scale out. In this evaluation, we have combined DATAIlegro's solution with the SQL Server, as the DATAIlegro product is now Microsoft's Madison project (the SQL Server MPP scale-out solution).

[✦ Return to Top](#)

Strengths

- Microsoft continues to offer value for the price paid, giving high value with a low total cost of ownership (TCO). The purchase of SQL Server 2008 Enterprise Edition includes SQL Server Analysis Services (SSAS), SQL Server Reporting Services (SSRS) and SQL Server Integration Services (SSIS), which means online analytical processing (OLAP), reporting and data integration for extraction, transformation and loading (ETL) are included in the entry price, although they normally require separate servers.
- SQL Server 2008 scales from small warehouses to medium ones without a great deal of effort, adding many new features such as star joins, data compression and policy-based management for the data warehouse workload. Microsoft's acquisition of DATAIlegro is one example of Microsoft's aggressive stance on data warehouse scaling. Microsoft has also released certified configurations for SQL Server to assist customers in the configuration of a data warehouse environment.
- Worldwide support from the vendor is extensive, including partners, value-added resellers, third-party software and tools, and the wide availability of the SQL Server skill base.
- Microsoft's acquisitions of Stratature and ProClarity in 2007, and of Zoomex and DATAIlegro in 2008, demonstrate Microsoft's commitment to data management and integration to complement its mostly utilitarian approach from the past. The acquisition of DATAIlegro has added skilled personnel in data warehousing — particularly in marketing large data warehouse solutions — as well as augmenting Microsoft's field sales and support staff.

[✦ Return to Top](#)

Cautions

- As reported by Gartner clients during inquiries, Microsoft's solution remains a more manual management environment with respect to the large data warehouse and its associated infrastructure than the solutions of many of its competitors. High availability, scalability in a cluster and managing large amounts of storage are a few of the examples requiring a higher degree of skill and manual support. Each new release (including SQL Server 2008) helps with new functionality (such as policy management for the mixed workload), but managing large data warehouse environments remains a challenge. Gartner believes that the DATAlegro acquisition will add value to this issue in the future.
- One of the advantages of the Microsoft solution is also a challenge. Data warehousing and BI are related, but they are not the same topic. BI often leverages data marts (warehouse dependent or independent) and, as a result, the ability to build an isolated data mart solution using SSIS, SSAS and SQL Server 2008 can result in reduced governance. The problem with flexible solutions is that they demand careful governance.
- SQL Server runs only on Windows Server and, therefore, lacks the portability of most of its competitors. Many IT organizations do not consider SQL Server an option, as they are not willing to run production DBMS infrastructure on Windows Server in the data center environment.
- The availability and pricing of the Microsoft-DATAlegro appliance offering (code-named Madison) that combines Windows Server, SQL Server 2008 and the DATAlegro solution acquisition have not been published, making it difficult for potential customers to plan for its use, and more difficult for them to wait for it.

[✦ Return to Top](#)

Netezza

Netezza, at the end of 2007 and throughout 2008, emerged from the role of a "disrupter" to become a serious competitor in the marketplace. It continues to leverage its hardware acceleration strategy with multilayered processing, but has introduced complex and large data set processing beyond the warehouse. Specifically, its work with ISP partners leverages the use of PowerPC chips in its architecture.

[✦ Return to Top](#)

Strengths

- Netezza has emerged from the shadows of opportunistic markets with a growing number of implementations in which its solution serves as the primary warehousing platform.
- A growing number of customers are using Netezza's solution as an enterprise solution. Netezza offers an appliance solution that largely eliminates the need to balance hardware and software implementations in the environment. Because of this, Netezza actually plays in two segments of the market — it is an add-on for existing warehouses as an appliance-based data mart and an enterprise warehouse for smaller implementations with less complex workloads.
- Netezza recently added the capability to store and process spatial data. Netezza's focus is not only on data warehousing, but on complex data sources. However, Netezza continues to evolve its DBMS to include emergent data types, instead of remaining solely an appliance for relational or more traditional data types.
- Netezza has a strong track record of new customers, with more than 200 customers at the end of January 2008.

[✦ Return to Top](#)

Cautions

- Netezza's prices are no longer a force for disruption. Other vendors have responded with published pricing and discounting, and have introduced entry-level solutions. Netezza prospects should no longer assume that pricing is an automatic "win" with this vendor.
- Netezza is very good at isolating POC constraints when competing head-to-head with other vendors. Prospects are advised that POC results, while excellent and valid, are often based on isolated workload situations (single workload type) or leverage Netezza's massive hardware strategy without demonstrating that its solution is subject to gradual performance reductions as data volumes and connected user counts increase.
- Netezza references report conflicting performance for a full-query queue. Some customers report performance efficiency so high that the query queue remains empty. But other customers report gradually falling performance as the query queue grows. Other Netezza references claim that query processing is completed so quickly that the queue never exceeds 20 queries.
- Netezza presents a conundrum. Its solution for hardware and software balancing is beneficial to implementation and overall management efforts for the warehouse environment. However, organizations that exceed the current Netezza appliance volume can face issues in upgrading to larger configurations. Older Netezza configurations will need to be replaced. In 2008, Netezza added the ability to add new racks, making the system field upgradable. It also offers capacity on demand, enabling the customer to purchase a larger system than needed, paying only for what storage it needs and then purchasing additional storage as required.

[✦ Return to Top](#)

Oracle

Oracle remains a leader in data warehousing, enjoying more than 48% of the RDBMS market (see "Market Share: Relational Database Management System Software by Operating System, Worldwide, 2007"). In 2008, Oracle added a data warehouse appliance offering — the HP Oracle Database Machine with the HP Oracle Exadata Storage Server — sold and serviced by Oracle.

[✦ Return to Top](#)

Strengths

- Worldwide support and customer experience make Oracle a solid choice for organizations seeking access to a wide experience base.
- Oracle Real Application Clusters (RAC) with Automatic Storage Management (ASM) is becoming accepted as an enterprise-level DBMS platform for data warehousing, capable of supporting large data warehouses (defined in the Market Definition section as those bigger than 20TB). 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 service-level agreements (SLAs) being implemented.
- Oracle now has four distinct data warehouse solutions: Oracle Database 11g (the DBMS stand-alone); Oracle Reference Configurations (certified server and storage configurations); Oracle Optimized Warehouse (off-the-shelf appliances from its hardware partners); and HP Oracle Database Machine (a data warehouse appliance with storage optimized for data warehouses [HP Oracle Exadata Storage Server] based on the Oracle Database 11g RAC, ASM and HP hardware, sold and serviced by Oracle). Oracle continues to extend the stack, now to hardware, giving customers a single

vendor for support.

- Oracle Database 11g has added enhanced materialized view and cube management (notably transparent SQL access and incremental update). This increases Oracle's capability to deploy end-user optimization layers with features not found in other DBMSs. It also includes enhancements to Oracle's partitioning option, including Partition Advisor, which suggests types of partitioning to enhance performance based on the database schema.
- Oracle's solution is the most portable data warehouse platform on the market — running on most hardware with Linux, Unix or Windows — and includes a free ETL tool (Oracle Warehouse Builder) with a Data Quality option.

[✦ Return to Top](#)

Cautions

- Gartner client inquiries indicate that Oracle's DBA team has a higher full-time equivalent (FTE) commitment — primarily from a higher degree of manual administration — than some other DBMS solutions for data warehousing. Conversely, in these same inquiries, many Oracle DBAs report that this perception is no more than a reaction to the high degree of control provided by the Oracle DBMS. Oracle has made progress in this area with automated management tools (for example, the Automatic Database Diagnostic Monitor [ADDM] and ASM), the release of the HP Oracle Exadata Storage Server and by providing auditing tools, advisor tools and other metadata-driven system analysis capability.
- Gartner clients continue to raise as an issue Oracle's pricing and contract practices. Recently, with price adjustments from Oracle and the current economic environment, we have seen an increase in Gartner inquiries related to pricing issues with Oracle (see "What Oracle's U.S. Price Increases Mean to Your Organization"). Another issue for customers is understanding which features of Oracle's solution are optional and priced separately (for example, Oracle Enterprise Manager Management Packs, partitioning and compression). Organizations are encouraged to remain diligent in assuring compliance with Oracle licensing (see "Prevent Oracle Technology License Compliance Issues").

[✦ Return to Top](#)

Sand Technology

Sand Technology is a small DBMS vendor that has had a column-store DBMS in the market for approximately eight years.

[✦ Return to Top](#)

Strengths

- Sand Technology has tried several different marketing approaches over the years, including branding the DBMS as an archive engine because of the high compression ratios achieved with column-storage DBMSs. In fact, because of its use of tokens in addition to the column store, Sand's solution achieves greater compression than other DBMSs.
- Sand has had a partnership with SAP since 2004 as a near-line data store (as opposed to an offline archive) for the SAP Business InformationWarehouse (BW) in installations where the size of the BW has grown to a degree that it is affecting performance. By integrating with NetWeaver (SAP's middleware), an SQL query to the BW can be routed automatically to either the BW or the Sand near-line storage engine. The performance degradation is minimal and the

transparency to the end user is an excellent feature.

- Sand continues to have a loyal client base. With new clients being slowly added, both from its partnership with SAP and from native Sand products, it will remain in the market as a viable vendor.

[✦ Return to Top](#)

Cautions

- Because of Sand's small size — it has approximately 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 especially true now that there are new entrants to the column-store DBMS space (Infobright, ParAccel and Vertica, as well as the mature Sybase IQ).
- The long-term partnership between SAP and Sand makes Sand a candidate for acquisition by SAP in the next five years. This would position Sand as a near-line storage engine specifically for the SAP BW, as well as integrating the Sand DBMS as an analytic engine (because of the analytic engine capabilities of a column-based DBMS, as described in the Market Overview section). This could be an issue for Sand clients not using its solution with SAP, as SAP tends to focus on its own products.
- Sand's solution, as a column-based DBMS, faces the challenge of proving performance in an enterprise data warehouse environment (see Market Overview section).

[✦ Return to Top](#)

Sun Microsystems-MySQL

MySQL was acquired by Sun Microsystems in January 2008 and is the most widely used open-source DBMS engine (see "Market Share: Relational Database Management System Software by Operating System, Worldwide, 2007").

[✦ Return to Top](#)

Strengths

- MySQL has continued to mature, especially since the acquisition by Sun Microsystems in 2008 and with the release of MySQL 5.1 — this has brought new functionality, growth of professional services, the addition of Sun's sales force and the continued addition of many new third-party software vendors. MySQL Enterprise, offering a complete installable version with tools to manage the installation and operating environment, continues to gain rapid market acceptance. Many clients are beginning to use MySQL as a data warehouse engine for small data warehouses of approximately 200GB to 500GB. Generally, many data warehouse implementations begin small and grow over time. MySQL can benefit from this and will see the same pattern, as its scalability is proven over time.
- MySQL has several references with mutiterabyte data warehouses in production using a technique Sun calls "sharding." This technique splits the database into smaller pieces of less than a terabyte. Although this requires more resources to manage the database and associated storage, it does represent another step in the direction of large data warehouse capabilities.
- Even after its acquisition by Sun, the MySQL solution still maintains a low price point — a free license with support subscriptions ranging from \$599 per year per server to \$40,000 per year for the unlimited server license of MySQL Enterprise.
- BrightHouse, Infobright's column-oriented engine, uses MySQL to create an analytic data warehouse solution. This is possible because

of the architecture of the MySQL DBMS, which allows the MySQL solution to work with multiple storage engines.

- Sun continues to enhance the management tools available as part of the MySQL Enterprise offering, recently adding the SQL Analyzer tool to assist in optimizing the performance of poorly executing SQL statements.

[✦ Return to Top](#)

Cautions

- MySQL continues to lack references for data warehousing that break the 1TB barrier in a single instance of the DBMS (see sharding in Sun-MySQL's Strengths section). It will need to demonstrate scaling above 10TB in a mixed workload to dispel the perception of a lack of scalability of the MySQL solution in a data warehouse. We believe this is one of the major development efforts under way as a direct result of the acquisition by Sun (see "Combining Sun and MySQL (Part 2): New DBMS Market Dynamics").
- Sun is facing increased competition from some of the new entrants using open-source DBMS technology, such as EnterpriseDB, which is just beginning to support data warehousing with EnterpriseDB GridSQL.
- MySQL still lacks many of the special features necessary to be a serious contender for large data warehouses. Although MySQL has some basic functionality for workload management (such as storing query statistics) and, with MySQL 5.1 has added partitioning, it will need to add more control and automatic management functionality to handle large data warehouses and the mixed workload.
- The low entry cost of using MySQL does not always equate to low TCO, as the cost to manage a large data warehouse without the broad availability of management tools (as with the larger, more mature data warehouse DBMSs) leads to the use of resources to perform these management tasks manually.

[✦ Return to Top](#)

Sybase

Sybase was the first of the column-store DBMS systems (the design is well adapted to analytics). The increasing number of column-store DBMS entries to the market and their appeal to end-user organizations appear to validate the design.

[✦ Return to Top](#)

Strengths

- Sybase recognized the market trend for data marts and departmental initiatives ahead of most of the market — as far back as 2006 — and capitalized on this trend. Sybase IQ Analytics Server already exhibits good scaling to enterprise capability, despite its marketing focus (especially in the financial market space). Gartner expects Sybase to expand its marketing to leverage the entry-level capability against enterprise scaling in 2009. Sybase has focused its sales force on the areas in which it stands out — Sybase IQ Analytics Server and Mobility. The company has seen revenue growth for Sybase IQ Analytics Server at times exceeding 80% quarter over quarter. This focus and rebranding has allowed Sybase to gain new customers, and increase the size of its implementations and deployments.
- Sybase IQ Analytics Server achieves data compression ranging from two times compression to five times compression, depending on the structure of the data. Because analytics typically makes use of fewer columns but larger numbers of rows, Sybase IQ Analytics Server

performs very well for analytic applications. The company has been consistently winning POCs with analytic applications, sometimes crushing the competition with performance 100 times greater. This makes Sybase IQ Analytics Server an extremely desirable DBMS platform for an analytic data mart to optimize and enhance an organization's overall data warehouse architecture.

- Sybase has a strong alliance with the IBM System p division. This channel has resulted in the Sybase IQ Analytic Data Warehouse Appliance, offered in three different scale factor configurations, based on the IBM System p platform, and sold and supported by a system integrator, mLogica. This is part of Sybase's overall strategy to build complete reference architectures for its target markets in collaboration with infrastructure and tools vendors. In addition, Sybase has implemented this strategy by adding its ETL tool to Sybase IQ Analytics Server and real-time loading solutions, expanding its capabilities and leading to additional prospects with a more comprehensive offering.

[✦ Return to Top](#)

Cautions

- Sybase has had to fight the perception of being smaller than IBM, Microsoft, Oracle and Teradata in the DBMS market. Although it has made great strides in changing market perception, it must work hard to prevent the previous perception from returning.
- Although Sybase IQ Analytics Server has a large installed base of approximately 1,500 customers, it faces strong competition from new column-based DBMS vendors, such as ParAccel and Vertica. Gartner believes the short-term effect of this competition will be to lengthen sales cycles and increase cost of sales due to a higher number of required POCs than in previous years.
- As Sybase expands into the enterprise data warehouse space, it will face increased competition from the incumbent vendors and more difficult POCs. While Sybase IQ Analytics Server remains ahead of the column-based newcomers in execution and has shown the ability to scale to EDW solutions, the vendor's challenge will be to respond to new market demands by offering a wider variety of data warehouse solutions, giving customers a clear vision for full-scale enterprise data warehouse solutions.

[✦ Return to Top](#)

Teradata

Teradata offers several data warehouse appliances combining hardware, operating system and DBMS. Its offerings include entry-level-priced solutions, data marts and data warehouses.

[✦ Return to Top](#)

Strengths

- Teradata allows customers to engage in several market microtrends. First, with the publication of lower pricing and the introduction of discounting, it has addressed pricing issues. Second, organizations are seeking entry-level solutions that exhibit mature market experience. With the introduction of the Data Mart Appliance 551, the Extreme Data Appliance 1550 and the Data Warehouse Appliance 2550, Teradata is promoting a solution that addresses the "learn as you grow" mentality of new data warehousing entrants.
- Teradata offers the opportunity for organizations to leverage 30 plus years of architecture specific to data warehousing, design and deployment experience.

Teradata also continues to show solid growth, with more than 19 quarters of revenue growth, specifically from data warehousing. Its broad customer base (approximately 1,000 customers) also contributes to its consulting expertise, including tools to track the development of warehouse design and deployment.

- Because of Teradata's architecture, it is well positioned to support the new, modern mixed workload, as proven with both its Active Data Warehouse and Dual-Active Data Warehouse. Also, its complete solutions (including the Teradata Data Warehouse), data models and professional services dedicated to data warehousing set it apart from the rest of the market. Finally, the SUSE Linux version has been in production for two years and takes advantage of the price competitiveness of Linux.
- Teradata's management software is a clear strength, as it manages the entire data warehouse environment, from the operating system to the workload, with software to manage the mixed workload, including a priority scheduling manager to prioritize the workload by application and/or groups of users.

[✦ Return to Top](#)

Cautions

- A concern for Teradata through 2011 remains the fierce competition from the mature DBMSs (IBM DB2, Microsoft and Oracle), as they become stronger in supporting a mixed workload in the sub-10TB range.
- Many new end-user organizations deploying new or revised data warehouses will implement them using a competitor based simply on enterprise standards, because they lack the experience base necessary to discern the more advanced requirements of mixed workloads, high availability and analytics optimization.
- Common data warehouse practices have seen a renewal in single-vendor appeal. Teradata offers only the data infrastructure for the BI environment. Other vendors that offer a complete stack — from extraction to reporting and OLAP — will renew claims that "Teradata is proprietary." This is a false claim, as the DBMS is no more proprietary than DB2, Oracle, SQL Server or Sybase IQ Analytics Server (among others) and, more important, the hardware (only Bynet is built to custom specifications). Organizations should ignore these claims and instead focus on the decision criteria regarding mixed workload demands, balanced system management and data optimization advantages, which are persistent and pervasive needs in the data warehouse DBMS market.

[✦ Return to Top](#)

Vertica

Vertica, the newest entrant to the data warehouse DBMS market, markets a column-store analytic DBMS with a number of functional capabilities added for high performance and high availability.

[✦ Return to Top](#)

Strengths

- Vertica's solution has had very strong early adoption as an analytic data mart, with more than 60 customers in less than a year. The DBMS is inexpensive, with a pricing model based on the amount of source system extracted data, not based on users, servers, chips or cores. Its fast adoption is also a result of its simple installation and broad portability across hardware systems. References report that they are able to set up Vertica data warehouses very rapidly,

sometimes in hours. This is partly because of a feature in the Vertica solution — automatic database design — that requires less optimization of the model.

- Vertica's DBMS has many features that help to set it apart from other DBMS engines, such as built-in high availability (including active replicas, auto-node recovery and no single point of failure shared-nothing architecture) and data compression (additional to the automatic compression realized as a column-store DBMS).
- Vertica's solution was the first DBMS with a "cloud" implementation, using Amazon Elastic Compute Cloud (EC2). This has been an advantage in POCs, as personnel from Vertica have access to the cloud rather than getting through customers' firewalls. Further, the implementation and set up is very fast, sometimes as little as an hour. This is also a strong offering for development purposes, as deployment is rapid and it can be discarded easily when finished.

[✦ Return to Top](#)

Cautions

- There are many entrants in the column-store DBMS space, which will make differentiation difficult. This clearly weighs toward more mature products with an installed base and makes it more difficult for newcomers such as Vertica.
- Vertica is lacking in large references running in the multiple terabyte range with large numbers of users, although this is partly because of the high degree of compression in a column-store DBMS.
- Vertica's solution, as a column-store DBMS with little mixed-workload management, is limited to data marts for analytic applications with a lower number of users. Gartner believes this will change in time, as new functionality is added to broaden its capabilities as a data warehouse DBMS.

[✦ Return to Top](#)

The Magic Quadrant is copyrighted 23 December 2008 by Gartner, Inc. and is reused with permission. The Magic Quadrant is a graphical representation of a marketplace at and for a specific time period. It depicts Gartner's analysis of how certain vendors measure against criteria for that marketplace, as defined by Gartner. Gartner does not endorse any vendor, product or service depicted in the Magic Quadrant, and does not advise technology users to select only those vendors placed in the "Leaders" quadrant. The Magic Quadrant is intended solely as a research tool, and is not meant to be a specific guide to action. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

© 2008 Gartner, Inc. and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner's research may discuss legal issues related to the information technology business, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The opinions expressed herein are subject to change without notice.