

(<https://www.gartner.com/home>)

LICENSED FOR  
DISTRIBUTION

# Magic Quadrant for Data Management Solutions for Analytics

**Published:** 13 February 2018    **ID:** G00326691

**Analyst(s):** Adam M. Ronthal, Roxane Edjlali, Rick Greenwald

## Summary

The data management solutions for analytics market is evolving as the cloud's position solidifies, use cases for Hadoop clarify, logical data warehouse adoption grows, and Chinese vendors expand abroad. Against this dynamic backdrop, this report will help you find the right vendor for your business.

## Market Definition/Description

We define a data management solution for analytics (DMSA) as a complete software system that supports and manages data in one or more file management systems (usually databases). DMSAs include specific optimizations to support analytical processing. This includes, but is not limited to, support for relational processing, nonrelational processing (such as graph processing), and machine learning and programming languages such as Python and R. Data is not necessarily stored in a relational structure, and multiple models can be used – for example, relational, XML, JSON, key-value, text, graph and geospatial.

Although the traditional data warehousing use case remains foundational to most organizations' analytics initiatives, there is also interest in the ability to manage and process increasingly diverse formats for both internal and external data. A complete DMSA must therefore be able to accommodate a diverse range of data types. These may include interaction and observational data – from Internet of Things (IoT) sensors, for example – as well as nonrelational data, such as text, image, audio and video data.

The breadth and scope of associated roles and skills is also expanding as organizations engage with new use cases that deliver a fuller understanding of data from an increasing number of sources.

We define four primary use cases for DMSAs that reflect this diversity of data and use cases (see also Note 1):

- Traditional data warehouse
- Real-time data warehouse
- Context-independent data warehouse
- Logical data warehouse (LDW)

Our definition also states that:

A DMSA is not a specific class or type of technology.

A DMSA may consist of many different technologies in combination. However, any offering or combination of offerings must, at its core, be able to provide access to data under management by open-access tools via standard APIs like Open Database Connectivity (ODBC), Java Database Connectivity (JDBC), representational state transfer (REST) and Object Linking and Embedding Database (OLEDB).

A DMSA must provide data availability to independent front-end application software, include mechanisms to isolate workload requirements, and control various parameters of end-user access within managed instances of data.

A DMSA must have administrative control of the data it is using. This means that it must control how data is persisted, accessed, governed and secured.

There are many different delivery models for DMSAs, such as stand-alone DBMS software, certified configurations or reference architectures, database platform as a service (dbPaaS) offerings and data warehouse appliances. These are evaluated together in our analysis of each vendor.

## Magic Quadrant

**Figure 1.** Magic Quadrant for Data Management Solutions for Analytics



Source: Gartner (February 2018)

### Vendor Strengths and Cautions

#### Action

Action (<https://www.action.com/>), which is headquartered in Palo Alto, California, U.S., offers the Action Vector analytics platform for analytical workloads, Action Vector in Hadoop, and Action X for combined operational and analytical processing. The Action Vector analytics platform can also be deployed on Amazon Web Services (AWS) and Microsoft Azure with a bring-your-own-license model, or via an Amazon Machine Image (AMI) for the community-supported free edition.

#### STRENGTHS

**Renewed investment in DMSA:** Action did not appear in the 2017 edition of this Magic Quadrant, due to a change in its strategy and roadmap. However, after bringing in new leadership, it is now reinvesting in its Vector technology to meet the demand for analytics.

**Performance:** Actian Vector is a column-oriented, in-memory DBMS that uses vector processing for query execution. Reference customers praised the technology's performance.

**Value for money:** Many reference customers praised Actian's price/performance. Actian scored better for value for money than in any other category in our survey of reference customers.

## CAUTIONS

**Cloud support:** Actian has yet to offer a strong cloud platform as a service (PaaS), even though the cloud is fast becoming a standard deployment option. This has limited Actian's ability to address its potential customer base. However, Actian's recent release of a community edition AML, together with its plans to offer a fully managed enterprise PaaS option for Vector on multiple cloud platforms in 2018, should address this need.

**Disruptions in 2016:** Actian scored in the bottom quartile for many criteria in our reference customer survey, including support, product capabilities and overall experience of doing business with a vendor. We believe this was due to the internal disruption Actian experienced in 2016 when it briefly discontinued support for DMSA platforms prior to a comprehensive change of leadership. Rebuilding mutual confidence and market acceptance will require dedicated effort from Actian, but we have seen early indications that this is already underway.

**Ease of implementation and system availability:** Reference customers gave Actian below-average ratings for ease of implementation and system availability. Recently, however, Actian has embarked on proactive engagements with existing customers to improve overall product performance and reliability, which should help.

## Alibaba Cloud

Alibaba Cloud (<https://www.alibabacloud.com/>) is a global cloud computing company headquartered in Hangzhou, China. It offers a wide variety of services, such as ApsaraDB for RDS (relational database service) for MySQL, SQL Server and PostgreSQL; HybridDB for PostgreSQL, based on the open-source Greenplum Database; AnalyticDB for OLAP analysis; MaxCompute for large data warehouse implementations; and E-MapReduce for Hadoop. In addition, Apsara Stack Agility provides an on-premises private cloud implementation.

## STRENGTHS

**Broad product portfolio:** Like Amazon, Alibaba Cloud started out as a cloud retailer, and its approach offers customers a wide choice of services.

**Hybrid deployment approach:** Although Alibaba Cloud offers primarily a public cloud service, it also offers the possibility of deploying its Apsara products on-premises. This is a differentiating approach, as many other cloud providers offer only public cloud deployment alternatives.

**Large deployments:** Half of Alibaba Cloud's reference customers indicated that they had deployments larger than 100TB, and some had analytical datasets running into petabytes.

## CAUTIONS

**Ease of implementation:** Alibaba Cloud's reference customers rated its ease of implementation below average, even though this is a criterion by which cloud PaaS providers generally score very well. We believe this is due to the complexities of combining multiple



cloud services and Alibaba Cloud's expectation that clients understand which service to use for which use case.

**Evaluations, contract negotiation and value for money:** Alibaba Cloud came bottom in our reference customer survey for evaluation and contract negotiation experience, and in the bottom quartile for value for money. We believe this is due to Alibaba Cloud's relative newness to this market, and we think that these scores will improve as the company gains market experience.

**Marketing and positioning:** The wide variety of technology offerings provided by Alibaba Cloud makes it difficult to align them clearly with particular use cases. Alibaba Cloud would benefit from clarifying its product positioning and recommended architectures.

## Amazon Web Services

Amazon Web Services (<https://aws.amazon.com/>) (AWS) is a wholly owned subsidiary of Amazon, which is based in Seattle, Washington, U.S. AWS offers Amazon Redshift, a data warehouse service in the cloud. Amazon Redshift includes Redshift Spectrum, a serverless, metered query engine that uses the same optimizer as Amazon Redshift, but queries data in both Amazon Simple Storage Service (S3) and Redshift's local storage; Amazon S3, a cloud object store; Amazon EMR, a managed Hadoop service; and Amazon Athena, a serverless, metered query engine for data residing in Amazon S3. Additionally, the recently announced Amazon Neptune provides graph capabilities.

### STRENGTHS

**Cloud dominance:** AWS is the dominant cloud vendor by a significant margin, with only Microsoft even close in terms of market share and presence. This dominance provides increased network effects for all its services, because the sources of data for a DMSA use case are more likely to reside in an AWS service than in a service from any other cloud vendor.

**Best-fit solution approach:** AWS offers several different services that can be used for different DMSA use cases. This approach enables customers to select only those capabilities needed for a defined use case, and the cloud nature of all the services helps to reduce the complexity of supporting multiple products.

**Pricing:** AWS is a leader in low-cost, pay-as-you-go pricing, with discounting for optional term commitments. In addition, AWS has a spot-pricing model for Amazon EMR that enables clients to obtain additional resources at significantly lower cost in response to bids for excess capacity from other clients. AWS has also recently introduced per-second billing and serverless options for some services, which improves flexibility and cost optimization opportunities.

### CAUTIONS

**No on-premises offerings:** AWS's DMSA offerings are available only in the cloud. Some clients will need a hybrid environment that supports both on-premises software and cloud-based services, as either a temporary or a long-term solution. AWS offers multiple services in the cloud, but no software for on-premises use, although a partnership with VMware aims to bridge the gap between on-premises and cloud deployments for some use cases.

**Evolving Redshift architecture:** Amazon Redshift is available in configurations that include fixed amounts of computing power and storage within a cluster, and Redshift Spectrum to scale computing resources to query data in Amazon S3. For the local storage configuration, customers cannot scale resources independently, and resizing or changing machine instance

types requires cluster reconfiguration, which can take several hours while data is redistributed and places the cluster in read-only mode for the duration. Redshift Spectrum, which works against data in Amazon S3, is not subject to this issue, but cannot deliver the same levels of absolute performance as Redshift in all use cases, and it has a different consumption-based pricing model.

**Integration complexity:** By offering purpose-built, independent services, AWS can reduce the complexity of each, but this can increase the work required to integrate the separate services. Integration requirements must be addressed by customers, so the proliferation of service instances may entail significant effort. AWS Glue provides managed extraction, transformation and loading (ETL) services and a data catalog, which should help.

## Cloudera

Cloudera (<http://www.cloudera.com/>), which is based in Palo Alto, California, U.S., offers the Cloudera Enterprise platform, versions of which include Cloudera Enterprise Data Hub, Cloudera Analytic DB (for business intelligence [BI] and SQL workloads based on Apache Impala), Cloudera Data Science & Engineering (for data processing and machine learning based on Apache Spark and Cloudera Data Science Workbench), and Cloudera Operational DB (for real-time data serving based on Apache HBase and Apache Kudu). Through its "shared data experience" (SDX) technologies, the platform provides unified security, governance and metadata management across these workloads, as well as across deployment environments. Cloudera's platform is available on-premises; across the major cloud environments (including native object store support for Amazon S3 and Azure Data Lake Store); and as a managed service under the Cloudera Altus brand.

### STRENGTHS

**Market presence:** Of all pure-play Hadoop distributions, Cloudera's is the most successful in this market, according to Gartner's published revenue numbers. It also has the highest percentage of further purchase intentions among Hadoop vendors in our reference customer survey. This finding is corroborated by Gartner's contract review data. Additionally, in the same survey, Cloudera was the third most frequently considered but not selected vendor, behind Oracle and IBM, which indicates that it is the most popular of the pure-play Hadoop vendors.

**Products and packaging:** Respondents to our survey praised Cloudera's features and packaging, saying that all the components work well together. Cloudera's value-added components, such as Cloudera Manager, Kudu and Impala, were frequently mentioned. By offering several versions of its core platform, including Analytic DB, Data Science & Engineering and Enterprise Data Hub, Cloudera offers customers more choice, to help them meet their use-case requirements.

**Value for money and support:** Reference customers gave Cloudera higher-than-average scores for value for money, and Cloudera's average scores for service and support were in the top quartile. Reference customers frequently praised the quality of its technical support.

### CAUTIONS

**Potential erosion of core Hadoop stack:** Cloudera, like other Hadoop distribution vendors, is being challenged as new processing alternatives (such as Apache Spark) and new storage options (such as Amazon S3 for cloud object storage) do not require a full Hadoop stack.

Cloudera has addressed this risk by adding Spark to its distribution in 2013, and by offering direct access to files stored in Amazon S3 in 2016 and in Microsoft Azure Data Lake Store in 2016.

**Traditional data warehouse performance concerns:** Respondents to our survey rated Cloudera's performance for traditional data warehouse workloads in the bottom quartile. Although still high, the percentage of customers who reported using another solution for traditional data warehouse operations was lower for Cloudera than for other Hadoop-based vendors.

**Cloud maturity:** Cloudera Altus, a dbPaaS offering, is relatively new, still gaining features and capabilities, and yet to gain significant traction with customers.

## GBase

GBase (<http://www.gbase8a.com/>) , a trading name of Tianjin Nanda General Data Technology, which is based in Beijing, China, offers GBase 8a, a relational massively parallel processing (MPP) data warehousing platform; GBase Infinidata 8a, a data warehouse appliance; GBase UP, an LDW platform supporting data virtualization between GBase 8a, Hadoop and other platforms; and GBase cloud DB (GBase 8a), available in the QingCloud app center.

### STRENGTHS

**Proven success in China and early traction in global markets:** GBase has registered strong results in China's very large market, with petabyte-scale deployments in the finance and telecom sectors. It is also gaining traction in the telecom sector in South America, Africa, Eastern Europe and North America.

**Customer loyalty:** GBase scored well in our survey for customer persistence. Slightly more than 90% of its reference customers stated that they intended to purchase additional licenses from the company within the next 12 months, and even more said they would recommend GBase to others.

**LDW vision:** With GBase UP, GBase aims to productize support for the LDW. Although it is too soon to assess this product's success, it shows early promise.

### CAUTIONS

**Branding:** GBase's branding of the GBase DBMS is very similar to Oracle's branding. This raises concerns about the long-term viability of the company's products in the global market – it will need to change if GBase is to continue to grow outside its home market of China.

**Customer satisfaction:** Surprisingly, given its strength in terms of customer persistence, GBase scored below average for customer satisfaction in our survey. Of the vendors in this Magic Quadrant, it received some of the lowest scores for pricing, contract flexibility and value for money, and its score for ease of implementation was in the bottom third.

**International presence:** GBase still sells predominantly in China. Additional investment will be required to make its products appeal to the global market, and especially to North America and Europe.

## Google

Google (<http://www.google.com/>) , based in Mountain View, California, U.S., is a wholly owned subsidiary of the Alphabet holding company. Google Cloud is the part of Google that focuses on delivering solutions and services to the business market. Google's dbPaaS offerings in the Google Cloud Platform include BigQuery, a managed data warehouse offering; Bigtable, a nonrelational wide-column DBMS; Cloud Dataproc, a managed Spark and Hadoop service; and Cloud Dataflow and Cloud Pub/Sub, both focused on real-time stream data processing.

## STRENGTHS

**Modern cloud architecture and pricing models:** Google has invested heavily in its cloud platform, and its products take advantage of a high-speed network and modern pricing options focused on serverless and performance-metric-based approaches. BigQuery and the other DMSA offerings and components (such as TensorFlow) have, like many Google Cloud Platform offerings, been deployed internally at Google for years.

**Ease of implementation and value for money:** Reference customers praised the Google Cloud Platform for its ease of implementation and value for money, saying that it allows them to focus on business problems rather than managing a data warehouse. Unsurprisingly, Google, along with other cloud vendors, received some of the best scores for rapid implementation.

**Global presence with strong focus:** Google is developing its enterprise offerings in the DMSA market and making good progress with some early, high-profile customer wins. As a global organization with abundant resources, we expect its DMSA offerings to mature rapidly.

## CAUTIONS

**Support:** Respondents to our survey expressed concerns about Google's support, which they scored in the bottom quartile. One respondent described the troubleshooting process as overly technical.

**Platform maturity and documentation:** The Google Cloud Platform still gives the impression of having a rapidly developing ecosystem. Product documentation and overall platform maturity need improvement.

**Developing partner ecosystem:** Reference customers reported that the Google Cloud Platform is still not a priority for third-party vendors and system integrators, which tend to focus on traditional on-premises DMSA platforms and other major cloud platforms. Google has, however, taken steps to develop its partner ecosystem, so we expect the situation to improve.

## Hortonworks

Hortonworks (<https://hortonworks.com/>) is based in Santa Clara, California, U.S. It offers a Hadoop distribution called Hortonworks Data Platform (HDP); Hortonworks DataFlow for streaming data delivery and ingestion; HDInsight Hadoop service for Microsoft Azure; and Hortonworks Data Cloud Hadoop service for Amazon AWS. In addition, the company recently introduced Hortonworks Dataplane Service, a unified architecture to manage, govern store, process and access datasets across multiple use cases.

## STRENGTHS

**Expanding capabilities and strong partnerships:** Hortonworks started out by supporting batch-oriented analytical processing for a wide variety of data. Its use cases have since expanded to support new demands, such as for cloud-based data lakes with access to Amazon S3 and IoT use cases with DataFlow. Additionally, Hortonworks has partnered with Microsoft and IBM to provide its capabilities to a wider range of customers.

**Rich open-source-based components:** Hortonworks remains a major contributor to the Hadoop community, but it also draws on contributions from other open-source initiatives, such as Apache Spark, Apache Zeppelin for agile analytics, and Apache Atlas for data governance and metadata management.

**Customer satisfaction:** A relatively large proportion of Hortonworks' reference customers would recommend its technology to others. They praised the vendor's open-source approach, but indicated that running the solution requires dedicated skills and alignment with specific use cases and business outcomes.

## CAUTIONS

**Strict adherence to open-source stack:** Although seen as an advantage by some, Hortonworks' strategy of adhering strictly to the open-source stack means it continually has to find new ways to differentiate itself.

**Value for money and product capabilities:** Hortonworks scored in the bottom quartile in our survey for both value for money and overall product capabilities. This is largely due to an adjustment of expectations, as our reference customer survey and other client interactions indicate that organizations now have a clearer understanding of which use cases are best addressed by Hadoop and, by extension, Hortonworks. Hortonworks therefore needs to focus on key use cases where its technology has had proven success.

**Ease of implementation:** Overall, reference customers gave Hortonworks a lower score for ease of implementation than they did other vendors in this Magic Quadrant. This probably indicates that they needed help in the form of professional services.

## Huawei

Huawei (<http://www.huawei.com/>), which is based in Shenzhen, China, offers the FusionInsight Big Data platform, a data management platform that combines components of Apache Hadoop, Spark and Storm, and FusionInsight LibrA, a proprietary MPP DBMS. Huawei has added industry-specific domain models, proprietary extensions to the Hadoop platform for event stream processing, graph and machine-learning capabilities, and a unified SQL engine that is compatible with its MPP database and runs on Hadoop. Additional enhancements have been made to the Hadoop scheduler with Huawei's Superior Scheduler, and to the supported Hadoop Distributed File System (HDFS) file formats with Apache CarbonData. Huawei's offerings are also available in the public cloud through partners.

## STRENGTHS

**Industry expertise:** Huawei has used its strength in the telecom industry to build comprehensive DMSA offerings that are broadly applicable to general use cases. Additionally, Huawei has rich experience in the fields of logistics, voice services and video enterprise intelligence, stemming from its deep understanding of the telecom sector.

**Customer loyalty:** Huawei's surveyed reference customers praised its technology and indicated strong intent to purchase additional licenses from the company during the next 12 months.

**Broad portfolio and capabilities:** The FusionInsight Big Data Platform blends Hadoop, stream computing and a relational MPP database in a single environment with a unified SQL engine. It can be deployed on-premises, in Huawei's public cloud, or in partners' public cloud environments. Extensions to the Hadoop open-source core have been made in the Apache

CarbonData distributed storage engine, HiGraph for graph processing, scheduling and orchestration, and the Elk interactive query language. Tight integration between the relational engine and Hadoop is also provided.

## CAUTIONS

**Market presence:** Outside Asia/Pacific, Huawei is not well-known in the data and analytics software market. Users of Gartner's client advisory service rarely ask about Huawei for DMSAs, and the surveyed reference customers of other vendors did not often consider Huawei. However, Huawei is expanding aggressively into international markets, with investments in local research-and-development facilities in Europe and North America.

**Inconsistent user experience:** Gartner's sources of information about Huawei's customer experience show inconsistent patterns that are difficult to resolve. We believe this is representative of inconsistency in terms of implementation and the overall experience of doing business with Huawei.

**Complex implementations:** Although Huawei's platform capabilities are broad, this breadth almost always entails integration effort as customers work to customize all or parts of the platform to meet their specific needs.

## IBM

IBM (<http://www.ibm.com/>), which is based in Armonk, New York, U.S., offers stand-alone DBMSs (Db2, Db2 for z/OS, Informix), appliances (PureData System for Analytics, PureData System for Operational Analytics, Integrated Analytics System, Db2 Analytics Accelerator), Hadoop solutions (BigInsights), managed data warehouse cloud services (Db2 Warehouse on Cloud), and private cloud data warehouse capabilities (Db2 Warehouse). IBM's BigSQL and Fluid Query provide a consolidated access tier to a wide range of DBMSs and Hadoop. In addition, IBM's DataFirst Method and Watson Data Platform support further evolution of hybrid cloud and on-premises deployment and management.

## STRENGTHS

**Unified DBMS engine:** Informix and the Netezza database engine (Netezza Platform Software) notwithstanding, IBM has spent much of the past few years moving toward a single DBMS engine to power its DMSA initiatives. The result is a renewed focus on Db2, which powers software-only, cloud and appliance form factors; it inherits some components of the Netezza-based PureData System for Analytics, which it will replace.

**Large and loyal installed base:** IBM is generally seen as a strong and strategic partner by its many customers. Over 94% of the IBM reference customers we surveyed intend to purchase additional licenses, products or features from the company over the coming 12 months. Additionally, IBM's reference customers would recommend its products and services to others, with 96% saying they would do so without reservation.

**Depth and breadth of portfolio:** As a large and mature organization, IBM has a deep product portfolio that spans a wide range of customer requirements, from product capabilities to deployment options.

## CAUTIONS

**Confusing branding and marketing strategy:** IBM's branding and marketing strategy confuses many customers. Reference customers highlighted an unclear strategic direction with the transition from Netezza-based to Db2-based offerings, and inadequate corporate

communication about the future of Informix. However, IBM's renewed attention on Db2, which is of strategic importance to the company, should help.

**Continued and late investment in declining markets:** IBM recently announced the Integrated Analytics System. This is a next-generation data warehousing appliance, based on Db2, that is Netezza-compatible and incorporates Spark, with a focus on expanding the data warehouse's role to include robust data science support as part of a centralized data hub. We believe, however, that the appliance market is in the middle of a decade-long decline. Furthermore, IBM's appliance offerings remain less flexible than those of some competitors, and have a poor track record for receiving regular hardware updates.

**Difficulty attracting new customers:** Although IBM remains a standard strategic partner for its existing customers, Gartner clients seeking advice about technology solutions for "greenfield" deployments rarely ask about IBM.

## MapR Technologies

MapR Technologies (<https://www.mapr.com/>), which is based in San Jose, California, U.S., offers its Converged Data Platform (CDP) in both open-source and commercial software editions. CDP features performance and storage optimizations using Network File System (NFS) and MapR-XD, a scalable POSIX-compliant data storage tier; MapR-DB, a nonrelational DBMS supporting key value, document, wide-column, graph and time series models; event-streaming capabilities (MapR-ES); high-availability improvements; and administrative and management tools. MapR Edge, a small-footprint edition of CDP, extends MapR's reach to edge-processing use cases common to IoT environments.

### STRENGTHS

**Converged platform:** MapR's CDP combines analytic capabilities with operational and real-time streaming. Reference customers like this "all in one" approach. MapR's reference customers reported using its product for all four DMSA use cases.

**Enterprise-readiness:** Reference customers consistently praised MapR's operational reliability, performance, flexibility and scalability. The Hadoop-compatible MapR filesystem (MapR-FS), nonrelational DBMS (MapR-DB), and ease of integration due to native NFS support were frequently identified as core differentiators.

**Growing portfolio and use-case coverage:** MapR has introduced robust streaming capabilities, scalable storage offerings and edge analytics capabilities over the past year as it has sought to expand its use-case coverage.

### CAUTIONS

**Market visibility:** MapR's market visibility remains limited. Gartner's client advisory service receives fewer inquiries about MapR than other vendors in this Magic Quadrant, and MapR rarely appears in contracts submitted to Gartner for review. MapR needs to keep increasing the market's awareness of its product and vision.

**Pace of ecosystem adoption:** Several reference customers noted that MapR's open-source packages can lag behind the latest versions available in the open-source community. Also, some complained that when new features were implemented, they did not adhere to the defined open-source standard. This is not entirely unexpected, however, as many vendors deliberately differentiate themselves from pure open-source offerings.

**Use-case selection:** Despite MapR's growing use-case coverage, reference customers reported that it remains important to select and validate appropriate use cases for MapR's platform.

## MarkLogic

MarkLogic (<http://www.marklogic.com/>) , which is based in San Carlos, California, U.S., offers a nonrelational multimodel DBMS, which it describes as "operational and transactional." The product is available in two editions: Essential Enterprise and a free Developer edition. Essential Enterprise can be deployed on-premises, in the cloud, and across hybrid infrastructures, including those of AWS, Microsoft (Azure) and Google (Google Cloud Platform), as well as on VMware, Pivotal (Cloud Foundry) and Red Hat platforms.

### STRENGTHS

**Robust multimodel features:** Although MarkLogic is a nonrelational document DBMS, it also supports atomicity, consistency, isolation and durability (ACID) transactions, and includes a search engine and a triple store to enable definition of semantic relationships using graph database capabilities. MarkLogic recently introduced an additional API to access rows, documents and triples in a single query.

**Focus on unified access to data:** MarkLogic positions itself as a unifying force for access to data. Essentially, it offers a new version of an operational data store. The advantage of this positioning is that organizations with complex data landscapes not only benefit from MarkLogic's approach and capabilities, but also gain increasing benefits as they combine more of their data and metadata in the operational data store.

**Customer satisfaction:** MarkLogic scored very well in the reference customer survey. It received the highest score for overall experience of doing business with a vendor and the third-highest score for vendor support. Users of Gartner's client inquiry service also reveal high levels of satisfaction with MarkLogic.

### CAUTIONS

**Restricted sales opportunities:** Large organizations with complex data landscapes are a natural fit for MarkLogic's capabilities. Smaller or less mature organizations may not recognize that investment in an additional data management platform would produce appropriate benefits.

**Market visibility:** MarkLogic has less visibility in this market than some other vendors in this Magic Quadrant. The number of inquiries received by Gartner that mention MarkLogic remains low, compared to this vendor's competitors, although inquiries about MarkLogic are increasing year by year.

**Change in focus:** MarkLogic shifted its focus to data integration two years ago, and it is still educating the market about the benefits of its vision. Although Gartner believes its vision to be sound, it remains to be seen whether MarkLogic can achieve widespread adoption.

## MemSQL

MemSQL (<https://www.memsql.com/>) , which is based in San Francisco, California, U.S., offers a distributed, scale-out SQL DBMS with an in-memory row store, along with a memory and disk-based column store that supports transaction and analytic use cases. MemSQL extends its DBMS platform by including real-time analytics with streaming data via Apache Spark or Apache



Kafka. MemSQL offers a free Developer Edition for nonproduction use and a paid-for Enterprise Edition that can be deployed on-premises or as a fully managed cloud service running on AWS or Microsoft Azure infrastructure.

## STRENGTHS

**Real-time low-latency capabilities:** MemSQL's combination of row and column store, with built-in streaming ingestion capabilities, represents a converged platform for low-latency analytics. Nearly half of MemSQL's reference customers report loading data continuously in near real-time — more than for any other vendor in this Magic Quadrant.

**Performance and ease of implementation:** Reference customers frequently praised the performance and ease of integration of MemSQL's product, as well as its operational overhead requirements. Its MySQL compatibility was also often identified as an advantage for ease of implementation.

**Hybrid cloud and multicloud vision:** Although some reference customers identified deployment issues with the managed cloud service, MemSQL's vision of a single platform, deployable in any cloud or on-premises, is strong.

## CAUTIONS

**Pricing concerns:** MemSQL is licensed by the amount of data under management. This prompted some reference customers to express concerns about pricing in environments where data volumes are growing, and about the long-term viability of this pricing model.

**Product maturity:** A relatively large proportion of MemSQL's reference customers identified absent or weak functionality, which indicates that its product is still maturing. However, MemSQL scored well in terms of performance, bugs and reliability, which indicates that when features are delivered, they are generally of high quality.

**Customer persistence:** MemSQL scored near the bottom of our survey results for customers' intentions to purchase additional licenses over the next 12 months. This finding is likely related to the pricing concerns mentioned above. MemSQL claims that its customers have, on average, nearly doubled their license capacity within the first 12 months, but could simply reflect initially small implementations that were rapidly outgrown.

## Micro Focus

Micro Focus (<https://www.microfocus.com/>), which is based in Newbury, U.K., offers the Vertica Analytics Platform. This platform is available as Vertica Enterprise, a columnar relational DBMS delivered as a software-only solution for on-premises use; Vertica in the Clouds, available as machine images from the AWS, Microsoft Azure and Google Cloud Platform marketplaces; and Vertica for SQL on Hadoop.

## STRENGTHS

**Unified engine that decouples storage and compute:** The Vertica DBMS engine is available in a traditional MPP scale-out version with direct-attached storage. It also supports major Hadoop distributions (from Cloudera, Hortonworks and MapR) via Vertica for SQL on Hadoop, which supports multiple file formats, including Apache Avro, Parquet and ORC.

**Performance, reliability and scalability:** Reference customers praised the Vertica Analytics Platform's performance, reliability and scalability. In addition, they frequently mentioned the product's maturity.

**Customer experience:** Micro Focus received above-average scores for almost every tracked indicator of customer satisfaction in our survey of reference customers. It received particularly high scores for product capabilities, price and pricing flexibility.

## CAUTIONS

**Acquisition of Hewlett Packard Enterprise (HPE) software:** There remains significant uncertainty in the market about Micro Focus' acquisition of HPE's software assets. Although there is every indication that Micro Focus intends to continue investing in the Vertica platform, and to keep supporting the Vertica brand strongly, potential customers may wish to exercise caution. They should expect some transitional friction until a long-term growth and investment strategy under Micro Focus is proven and yielding positive results.

**Cloud maturity:** Although the Vertica Analytics Platform is available in major cloud environments, reference customers reported that it feels much more like infrastructure as a service than a fully managed PaaS. Vertica 9 with Eon Mode pricing should alleviate some of these concerns.

**Maintenance and administrative concerns:** Reference customers reported challenges with adding and removing nodes, and poor performance for backup and restore operations. Administrative concerns were also identified in the prior Magic Quadrant, but recent product enhancements should improve manageability, backup and elasticity.

## Microsoft

Microsoft (<https://www.microsoft.com/>), which is based in Redmond, Washington, U.S., offers SQL Server as a software-only solution with certified configurations. It also offers the Analytics Platform System, an MPP data warehouse appliance. In addition, it sells Azure SQL Data Warehouse (a fully managed, MPP cloud data warehouse), Azure HDInsight (a Hadoop distribution based on Hortonworks), Azure Databricks (an Apache Spark-based analytics platform) and Azure Data Lake (a big data store and analytics platform) as cloud services.

## STRENGTHS

**Cloud focus and features:** Microsoft demonstrates a strong understanding of the market's cloud needs and required capabilities. With Azure SQL Data Warehouse, it addresses the growing interest in cloud data warehousing, but also hybrid on-premises and cloud use cases, and it is starting to demonstrate hybrid capabilities with stretch tables. Microsoft's Azure Data Catalog provides an integrated source of information across multiple cloud repositories. Additionally, Azure SQL Data Warehouse offers good support for LDW architecture, along with support for seamless integration with Hadoop and Azure Data Lake via Microsoft's PolyBase software.

**Value for money:** Respondents to our reference customer survey were generally pleased with the value for money they received from Microsoft. The vendor's relatively straightforward packaging and the absence of many expensive options appear to resonate well with its customers.

**Customer loyalty:** A relatively large proportion of Microsoft's reference customer respondents indicated that they would purchase more products or services from Microsoft in the next 12 months.

## CAUTIONS

**Pricing concerns:** Reference customers expressed concern about Microsoft's core-based licensing model for SQL Server, noting that as hardware vendors release CPUs with more cores, costs rise. Gartner, however, does not see this as a significant concern, as core-based licensing is an established industry practice and Microsoft also offers cloud pricing models that are not core-based.

**Training and support:** Microsoft scored near the bottom in our reference customer survey for both quality of end-user training and customer support. Microsoft has, however, indicated that it is investing heavily in documentation, training and support, and seeing positive results from these investments. We expect these scores to improve.

**Size of implementations and missing enterprise features:** Microsoft customers identified relatively low average sizes for their largest DMSA instance and the dataset regularly used for queries, which indicates that Microsoft is not yet fully meeting the needs of large enterprises in production environments. Additionally, they expressed the view that Microsoft's Always On feature in SQL Server for read scale-out is inferior to some competitors' offerings that offer read/write scale-out.

## Neo4j

Neo4j (<https://neo4j.com/>), which is based in San Mateo, California, U.S. and Malmö, Sweden, provides a graph platform that includes the Neo4j native graph database, graph analytics, the Cypher graph query language, data integration, and graph visualization and discovery tools. The company offers the open-source Neo4j Community Edition; Neo4j Desktop, which is free for developers and data scientists; and the paid-for Neo4j Enterprise Edition for production deployments.

### STRENGTHS

**Graph DBMS:** Neo4j is the leader in the stand-alone graph DBMS market in terms of both mind share and customer adoption (it claims over 10 million downloads of its Neo4j database software). Although many vendors are adding graph capabilities to their products, customers who need very robust graph capabilities with broad features, including analytic algorithms and tools to migrate from relational DBMS offerings, may find Neo4j's better suited to their needs.

**Ease of implementation, performance and training:** None of the reference customers we surveyed reported any issues with the implementation and use of Neo4j's platform. In addition, there were very few reports of product performance problems. Furthermore, Neo4j's reference customers scored the quality of its training higher than the average.

**Customer advocacy and support:** Neo4j received one of the top scores in the survey for customers' willingness to recommend its product to others. Additionally, respondents to our survey regularly praised Neo4j's support organization, with reports of quick responses and fast issue resolution.

### CAUTIONS

**Focus on graph solution:** Many larger and more mature vendors are adding graph capabilities to their products. Although their capabilities may not yet be as robust as Neo4j's in terms of features and performance, these vendors are likely to satisfy customers' initial need for graph capabilities, and may enhance their graph capabilities in response to increasing demand. Neo4j's sponsorship of openCypher — an open-source version of the Cypher query language — aims to broaden Neo4j's adoption as the graph standard, but it is still early days.

**Low purchase intentions and difficulty of doing business:** Neo4j had one of the lowest proportion of reference customers expressing intention to purchase more in the next 12 months, which is surprising given their strong advocacy. Surveyed customers scored Neo4j poorly for a variety of business-related measures, including licensing, pricing and pricing flexibility. Neo4j has recently modified its licensing terms and expects to offer cloud-hosted services in 2018.

**Upgrades and enterprise maturity:** Survey respondents reported issues with upgrading Neo4j's platform – they referred to challenging upgrades that sometimes had unexpected side effects. Additionally, some reported a lack of maturity or capabilities in several key enterprise areas, including security, monitoring and automation. The latest version of Neo4j's platform, which includes server-to-server encryption, should address at least some of these concerns.

## Oracle

Oracle (<https://www.oracle.com/>) is based in Redwood Shores, California, U.S., provides Oracle Database 12c, Oracle Exadata Database Machine, Oracle Big Data Appliance, Oracle Big Data Management System, Oracle Big Data SQL and Oracle Big Data Connectors. In addition, the Oracle Cloud service provides Oracle Database Cloud Service, Oracle Database Cloud Exadata Service and Oracle Big Data Cloud Service – a lineup to which Oracle Autonomous Data Warehouse Cloud will be added. Oracle's cloud portfolio also includes on-premises solutions in the form of Oracle Database Exadata Cloud at Customer and Oracle Big Data Cloud at Customer.

### STRENGTHS

**Technical vision and capabilities:** Oracle has been a leader in DBMS technologies for decades, and its capabilities make it one of the most prominent vendors in the DMSA market. The addition of Autonomous Data Warehouse Cloud should enable Oracle to meet the market's expectations for a cloud-based service. Oracle received the highest score in the reference customer survey for product capabilities.

**Integration with Hadoop distributions:** With Big Data SQL, Oracle extends its reach to a number of Hadoop distributions. It provides not only virtual access to data in Hadoop, but also advanced features such as predicate pushdown to targeted platforms.

**Market presence:** Gartner's data sources indicate that over 70% of DMSA purchase decisions are with customers' established DBMS vendors. This, coupled with Oracle's strong technical capabilities in the DMSA area, makes Oracle a popular choice for the use cases covered in this Magic Quadrant. Over 80% of Oracle's surveyed reference customers also used Oracle for their transactional systems.

### CAUTIONS

**Customer dissatisfaction:** Oracle charges premium prices and negotiates tough contracts, which can result in dissatisfied customers. A significant number of the Oracle customers who use Gartner's inquiry service are less than happy with its business practices.

**Unproven cloud capabilities:** Oracle's most viable cloud offerings are new and yet to prove their capabilities and licensing and pricing models in production environments. Oracle claims significant interest from existing enterprise customers, however, and is investing heavily in making its cloud offerings successful.

**Pricing and support:** Respondents to our reference customer survey frequently expressed concern about Oracle's pricing and support. The option-based approach to licensing is generally disliked by customers, especially for features that are seen to be necessary for effective production use. Oracle has, however, recently introduced Universal Credits for its cloud services and lowered database cloud prices for services in its cloud, which should help to alleviate some cloud-pricing concerns.

## Pivotal

Pivotal (<https://pivotal.io/>), which is based in San Francisco, California, U.S., offers the Pivotal Greenplum Database, an open-source MPP database based on PostgreSQL. It is available both as software and in the cloud on either AWS or Microsoft Azure infrastructure. Pivotal also offers an appliance in the form of the Dell EMC Data Computing Appliance, as well as GemFire, an in-memory data grid product. In addition, Pivotal sells a caching service called Pivotal Cloud Cache, which is based on GemFire and runs on the Pivotal Cloud Foundry platform.

### STRENGTHS

**Open-source alignment:** Pivotal made the Greenplum Database open-source in 2015, and the momentum it gained in this area in 2016 has continued, as Greenplum 5 is now compatible with PostgreSQL 8.4 and the company has begun to work on compatibility with PostgreSQL 9.x. In addition, Alibaba Cloud uses the Greenplum Database as the basis for some of its data warehouse offerings.

**Implementation for complex initiatives:** Pivotal's reference customers did not report any problems with complex implementations. This is at least partly because the Greenplum Database tends to appeal to technically minded users, who cope well with large, complex projects.

**Robust in-database analytic capabilities:** The Pivotal Greenplum Database incorporates Apache MADlib analytic libraries to provide a comprehensive suite of in-database analytic features "out of the box." These include mathematical, statistical, machine-learning, geospatial, time series, text analytics and graph algorithms. This makes the Greenplum Database well-suited to exploratory and data science use cases, in addition to traditional data warehousing.

### CAUTIONS

**Difficulty of implementation:** The Pivotal Greenplum Database appeals to end users who like to engage with a product on a very technical level. Pivotal's scores for ease of implementation were in the bottom third. Notably, Pivotal does not offer a REST API, which is increasingly a standard feature of products in this market.

**Software bugs and quality:** Pivotal's reference customers reported a relatively high number of issues with bugs. They also noted, however, that Pivotal's software quality has improved greatly in recent releases.

**Experience of doing business:** In our reference customer survey, Pivotal scored in the bottom quartile for experience doing business with a vendor. We believe this is at least partly due to the relative complexity of Greenplum Database implementations.

## Qubole

Qubole (<https://www.qubole.com/>) is based in Santa Clara, California, U.S. It offers the Qubole Data Service Enterprise Edition (QDS), a cloud-based Hadoop and Spark processing engine for data under its own management or managed by cloud object storage or other data management solutions, such as relational DBMSs. Qubole also offers Qubole Cloud Agents, which are add-on services to optimize cloud resource consumption and automate workloads. Qubole's offerings are available on AWS, Microsoft Azure and Oracle Cloud Platform.

## STRENGTHS

**Ease of implementation and experience of doing business:** Reference customers for Qubole praised QDS's ease of implementation, scoring it in the top-third of the vendors covered by our survey. As it offers a true PaaS, Qubole is differentiated from other Hadoop-based distribution vendors. Additionally, Qubole scored near the top in our survey for overall experience of doing business with a vendor. It came top for support experience.

**Storage-agnostic:** Qubole offers a PaaS that can manage data stored in multiple repositories, such as object stores and relational DBMSs. This is of particular value as the perceived requirement to colocate data for analytics is increasingly challenging organizations. In parallel, demand for flexibility and agility of access to data is increasing.

**Optimized use of cloud resources:** Qubole offers optimizations to reduce clients' cost of ownership, such as cluster autotermination and use of spot instances. Optimized processing is made possible by the separation of storage and computing resources — another factor providing strong differentiation from other Hadoop-based offerings.

## CAUTIONS

**Product capabilities and use-case alignment:** Qubole uses open-source processing frameworks such as Apache Spark, Hadoop supporting Presto, Hive and Pig, which are not suitable for all types of workload. Qubole scored in the bottom quartile for product capabilities in our reference customer survey. Enterprises considering Qubole should ensure their use cases match its processing capabilities.

**Enterprise-grade security capabilities:** Although Qubole offers security features such as encryption and Security Assertion Markup Language (SAML) support, enterprises should ensure that Qubole meets their security needs, such as for row- and column-based access control. This may be difficult, as the data may not always be under Qubole's management.

**Increased competition:** Qubole offers unique cloud deployment and processing flexibility. However, its position will be challenged by cloud query services that also allow direct querying of data, such as Amazon Athena and Amazon Redshift Spectrum. Many of these services, though not as flexible as Qubole's offering, may be good enough for certain uses.

## SAP

SAP (<http://www.sap.com/>) is based in Walldorf, Germany. It offers SAP HANA, an in-memory column-store DBMS that supports operational and analytical use cases. There is also SAP BW/4HANA, a packaged data warehouse solution. Both are offered as cloud solutions (for deployment in public and private clouds, and on SAP Cloud Platform) and as an appliance-like hardware reference architecture. SAP also offers SAP Cloud Platform Big Data Services, a cloud-based Hadoop distribution, and SAP Vora for Spark and Hadoop processing.

## STRENGTHS

**Performance:** SAP reference customers praised SAP HANA's performance. This approval is reinforced by Gartner's interactions with clients, which indicate that increased performance for analytical use cases is a significant driver of SAP HANA adoption.

**Maturity:** SAP reference customers expressed satisfaction with the maturity of its products. This view is reinforced by the finding that half of SAP's reference customers claimed deployment times of up to only three months.

**Satisfied reference customers:** Over 90% of SAP's surveyed reference customers indicated an intention to purchase additional DMSA product licenses from SAP in the next 12 months. This finding is underlined by the relatively large proportion of SAP reference customers who were willing to recommend the product to others.

## CAUTIONS

**Penetration beyond SAP's installed base:** SAP HANA is used primarily in support of SAP applications within the SAP ecosystem. Penetrating the larger, non-SAP market remains a challenge for SAP. This is partly due to pricing differences between the SAP runtime and enterprise use licenses (used for non-SAP applications), which indicates that SAP's focus is still very much on the SAP application ecosystem.

**Pricing and pricing flexibility:** Some respondents to our reference customer survey expressed concerns about the complexity of SAP's pricing and about its pricing flexibility. This, however, probably reflects a lack of understanding about the full use of SAP HANA, as opposed to runtime licensing.

**Data lake and cloud roadmap:** SAP's roadmap for data lakes on cloud object storage and for the separation of storage and computing resources is lagging behind what the market demands. SAP's offerings remain focused on Hadoop distributions and SAP Cloud Platform Big Data Services.

## Snowflake

Snowflake (<https://www.snowflake.net/>), which is based in San Mateo, California, U.S., offers a fully managed data warehouse as a service on AWS infrastructure. It supports ACID-compliant relational processing, as well as native support for document store formats such as JSON, Avro and XML. A native Apache Spark connector, R integration, support for user-defined functions, dynamic elasticity, temporal support and recently announced data-sharing capabilities round out the core offering. Snowflake is currently available only in the AWS cloud.

## STRENGTHS

**Customer satisfaction and ease of deployment:** Snowflake customers' satisfaction with its ease of implementation, flexibility of deployment and performance with very large datasets is strong. Snowflake's scores for ease of implementation and overall experience led those of all the other vendors covered by our reference customer survey.

**Market momentum:** Snowflake, while still a small vendor, is gaining market traction. This is apparent from strong growth in its customer numbers and steady growth in its overall ecosystem, as measured by the involvement of, for example, analytics and BI vendors, data integration vendors and service providers.

**Data sharing:** In addition to enabling organizations to implement new analytical use cases easily, Snowflake has introduced data-sharing capabilities that enable customers to easily share and monetize data stored in Snowflake.

## CAUTIONS

**AWS-only support:** Although AWS is the dominant cloud provider, the number of data warehouse solutions delivered as PaaS is increasing across a range of cloud providers, including Microsoft (Azure) and Google. Snowflake does, however, plan to extend its offering to other cloud providers soon.

**Requirement to load data into Snowflake:** To benefit from the service's processing flexibility and separation of storage and computing resources, data must first be loaded into Snowflake. Reference customers indicated that it is essential to match use cases to Snowflake's processing capabilities in advance. One customer noted that data loading was the most expensive aspect of an implementation. Another wished to be able to analyze data directly in Amazon S3, instead of having to load it into Snowflake.

**Pricing predictability:** Although Snowflake's score for pricing suitability was in the top third in our survey results, some respondents noted that it was difficult to keep costs under control. The recently announced move to per-second billing should help, however.

## Teradata

Teradata (<http://www.teradata.com/>) is based in San Diego, California, U.S. Its offerings include business and analytic consulting services, the Teradata Analytics Platform built on the Teradata Database, a software-only DBMS solution; Teradata IntelliFlex and IntelliBase appliances, and a range of cloud data warehouse solutions (all with MPP). Teradata IntelliCloud is an "as a service" cloud offering available on public cloud infrastructure (AWS and Microsoft Azure) and the Teradata Cloud (optimized infrastructure). Support for the LDW comes in the form of Teradata's Unified Data Architecture (UDA). Teradata QueryGrid (part of the UDA) provides multisystem query support via Teradata's own software, as well as via open-source Presto. Teradata also offers Aster Analytics and Hadoop support for Cloudera, Hortonworks and MapR distributions.

## STRENGTHS

**Adaptation to new demands:** Teradata has broadened its offerings in an effort to change the perception that it is suitable only for advanced, large-scale data warehouses. In response to demand, it now also offers free developer licenses and competitive base license pricing across on-premises and cloud deployment options.

**Product innovation:** Teradata innovates to meet the market's expectations. It has, for example, introduced health monitoring and cost-based optimization across the UDA. The Teradata Analytics Platform is innovative in that it enables multiplatform query and analytical processing, SQL, statistics and machine learning.

**Deployment and licensing flexibility:** With Teradata Everywhere, Teradata now supports all deployment models, with public cloud, Teradata Cloud and "as a service" IntelliCloud offerings, and on-premises IntelliFlex and VMware offerings. It supports hybrid – on-premises and cloud – deployments, and offers a single portable license model for all deployments.

## CAUTIONS



**Pricing and value for money:** Surveyed reference customers voiced concerns about Teradata's pricing and value for money. Some remarked that its pricing is complex, perhaps because of the company's shift to subscription-based pricing and packaging in 2017. We believe this is at least partly due to unrealistic perceptions of the true cost of subscription pricing and unrealistic expectations that the shift would translate into lower costs over the long term.

**Account management and customer satisfaction:** Teradata's scores from surveyed reference customers for account management and customer satisfaction were uncharacteristically below-average this year. We believe this reflects changing market demands, which are prompting customers to investigate alternative vendors (even if they do not actually leave Teradata), and that it may be a short-term phenomenon.

**Sales and marketing execution:** Teradata has made important changes to address the market's declining interest in data warehouse appliances, by, for example, introducing cloud and software-only offerings. However, Gartner's interactions with clients reveal that they still often fail to identify Teradata as a cloud data warehouse provider. Additionally, the change to a subscription pricing model has affected Teradata's revenue, if probably only temporarily as the company completes the transition.

## Treasure Data

Treasure Data (<https://www.treasuredata.com/>) is based in Mountain View, California, U.S. It provides the Customer Data Platform, a fully managed DMSA running on AWS infrastructure with availability regions in the U.S. and Japan. The Customer Data Platform provides a cloud data lake, combined with relational data marts. The abilities to ingest data from a wide range of sources, and to feed data to downstream data management platforms, are a focus for Treasure Data.

### STRENGTHS

**Data integration and IoT focus:** Treasure Data has focused on building flexible data integration capabilities, which are important for a data lake approach. Crowdsourced data integration via Fluentd and Embulk plugins provide streaming and batch-based capabilities. Treasure Data's early focus on IoT data provided a basis for much of its platform development.

**Advanced analytics capabilities:** Treasure Data supports native user-defined functions, built-in analytics capabilities for geospatial and statistical models, and strong integration with R and Python. Apache Spark integration is in beta testing.

**Support and pricing:** Respondents to our reference customer survey praised Treasure Data's support and pricing models. Treasure Data came third for overall support, and second for pricing flexibility.

### CAUTIONS

**Product maturity and learning curve:** Surveyed reference customers reported that Treasure Data's platform works well, but requires users with strong technical skills and capabilities to derive the most value from it. The company plans to improve these capabilities in 2018.

**Focus on data lakes:** Although Treasure Data does feed traditional relational data marts and enterprise data warehouses downstream, its platform's core strength remains its Hive- and Presto-based data lake capabilities. Respondents to our reference customer survey reported

regularly using Treasure Data's platform alongside other data warehousing solutions. The company can deliver relational cloud data marts, and claims that about one-third of its customers use its product in this way.

**Limited global presence:** Treasure Data's customers are almost all in Japan and the U.S., although they are often large enterprises with global presence. Customers in other countries may find localization and support resources limited.

## Vendors Added and Dropped

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

### Added

Actian

Alibaba Cloud

GBase

Micro Focus

Neo4j

Qubole

Treasure Data

### Dropped

1010data, which did not demonstrate that at least 10% of its customers were outside its home region, and therefore failed to meet the inclusion criterion requiring production customers from at least two geographic regions.

EnterpriseDB, which did not demonstrate that it fully supported at least two of the four defined use cases.

Hewlett Packard Enterprise, which completed the spin-off to, and merger of its software division with, Micro Focus in September 2016. This included the Vertica software product, which is now considered as part of this Magic Quadrant's evaluation of Micro Focus.

MongoDB, which did not demonstrate that it fully supported at least two of the four defined use cases.

Transwarp Technology, which did not demonstrate that at least 10% of its customers were outside its home region, and therefore failed to meet the inclusion criterion requiring production customers from at least two geographic regions.

## Inclusion and Exclusion Criteria

To be included in this Magic Quadrant, vendors had to meet the following criteria:

Each vendor had to have DMSA software generally available for licensing or supported for download for approximately one year (since 1 December 2016). We did not consider beta releases.

We used the most recent software release to evaluate each vendor's current technical capabilities. All software versions currently used in production environments were considered during direct interactions with vendors' reference customers, including the formal survey. For older versions, we considered whether later releases may have addressed reported issues, but also the rate at which customers refuse to move to newer versions.

Product evaluations included technical capabilities, features and functionality present in the product(s) or supported for download on 1 December 2017. Capabilities, product features and functionality released after this date could be evaluated at Gartner's discretion and in a manner deemed by Gartner to be appropriate to ensure that this Magic Quadrant's quality for Gartner's nonvendor clients. We also considered how such later releases might impact the end-user experience.

Each vendor had to identify 30 verifiably revenue-generating production implementations of its DMSA(s) at distinct organizations.

Each vendor must have had *either* a minimum of \$10 million (U.S. dollars) in revenue with at least 50% growth from calendar year 2015 to calendar year 2016, *or* more than \$40 million in revenue during the same period.

Revenue could be from licenses, support and/or maintenance. Revenue requirements for this Magic Quadrant are unchanged from those of the previous Magic Quadrant. We considered public sources of revenue reporting where possible, but if such sources were not available, and if Gartner market share data was insufficient or disagreed with what the vendor reported, we accepted a signed attestation from a senior executive of the vendor's finance department or its CEO.

Each vendor's production customer base had to include customers from at least three of the industry sectors listed in Note 2.

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

Each vendor had to demonstrate production customers from at least two of the geographic regions listed in Note 3. In addition, at least 10% of the verified production customer base must be outside the vendor's home geographic region.

Each vendor had to offer support for the evaluated DMSA product(s). We also considered products from vendors that control or contribute specific technology components to the engineering of open-source DBMSs and their support.

Any acquired product must have been acquired and offered by the acquiring vendor on or before 30 June 2017 for it to be included in the main evaluation of that vendor. Any later acquisitions would be represented by a separate dot until the next Magic Quadrant.

Each vendor had to offer significant value-added capabilities beyond simply providing an interface to data stored in other sources. We included in our assessments the capability of vendors to coordinate data management and processing from additional sources beyond the

evaluated DMSA.

Each vendor must support for at least two of the four primary use cases (for which, see Note 1). (T this is a change from last year's criteria, which required support for only one use case).

Each vendor must provide at least relational processing. Depth of processing capability and variety of analytical processing options are considered advantageous in the evaluation criteria.

Each vendor had to demonstrate an ability to deliver the services needed to support a DMSA through the establishment and delivery of support processes, professional services and/or committed resources and budget.

Products that exclusively support an integrated front-end tool that reads only from the paired data management system did not qualify for assessment in this Magic Quadrant.

We also considered the following factors when deciding whether products were eligible for inclusion:

Whether, and to what extent, they support relational DBMSs.

Whether, and to what extent, they support nonrelational DBMSs.

Whether, and to what extent, they support Hadoop distributions.

Whether, and to what extent, they support open-source solutions.

Whether multiple solutions are used in combination to form a DMSA; we considered this an eligible approach, provided the maturity and customer adoption of each solution could be demonstrated.

Deployment options: We considered cloud solutions viable alternatives to on-premises solutions; the ability to manage hybrid on-premises-and-cloud solutions increased vendors' chances of inclusion.

No specific rating advantage was associated with the type of data store used (for example, relational DBMS, graph DBMS, HDFS, key value DBMS, document DBMS, wide-column DBMS).

The following technology categories were excluded:

BI and analytical solutions that offer only an embedded DMSA or that embed a DMSA from another provider.

BI and analytical solutions that offer only a DMSA limited to the vendor's own BI and analytical solutions or whose customers only use the solution within the same vendor's stack.

In-memory data grids.

Hybrid transactional/analytical processing (HTAP)-only solutions.

Prerelational DBMSs.

Object-oriented DBMSs.

Gartner analysts were the sole arbiters of which vendors and products were evaluated in this Magic Quadrant.

## Evaluation Criteria

## Ability to Execute

The Ability to Execute criteria are primarily concerned with vendors' capabilities and maturity. These criteria also consider products' portability and ability to scale and run in different operating environments, which gives the customer a range of options.

Ability to Execute criteria are critical to customers' satisfaction and success with a product, so interviews with and survey responses from reference customers are weighted heavily throughout.

**Product or Service:** This criterion assesses vendors in light of increasingly divergent market demands — for traditional, logical data warehousing, real-time data warehousing, context-independent data warehousing and HTAP approaches to data management for analytics. The largest and most traditional portion of the analytics and data warehouse market is still dominated by demand to support relational analytical queries over normalized and dimensional models (including simple trend lines through complex dimensional models). DMSAs are increasingly expected to include repositories, semantic data access (such as federation/virtualization) and distributed processing in combination — as LDWs. All the traditional demands of the data warehouse remain. The real-time data warehouse use case also exhibits traditional requirements, plus the ability to accommodate streaming data, real-time data loading and real-time analytics support. Users expect solutions to become self-tuning, to reduce the number of staff required to optimize the data warehouse, especially as mixed workloads increase. Context-independent warehouses (CIWs) do not necessarily support mixed workloads (although they can), nor do they require the same level of mission-critical support. CIWs serve more in the role of data discovery support, data science initiative or "sandbox." CIWs are expected to meet the demands of ad hoc queries and varied processing options such as Python, ML, R and graph. In-database analytic capabilities and the ability to move analytic models from CIWs to traditional or operational models are also considered, as is flexibility of deployment (cloud, on-premises or hybrid).

**Overall Viability:** This criterion covers corporate aspects such as the skills of personnel, financial stability, investment in research and development (R&D), overall management of the organization, and the expected persistence of a technology during merger and acquisition activity. It also covers a company's ability to survive market difficulties. Vendors are also evaluated on their capability to establish dominance in meeting one or more discrete market demands.

**Sales Execution/Pricing:** This criterion examines the price/performance and pricing models of a vendor's DMSA(s), and the ability of its sales force to manage accounts (judged by feedback from our clients and the survey of reference customers). It also considers the market share held by the vendor's DBMS software. Also evaluated is the diversity and innovativeness of packaging and pricing models, including the ability to promote, sell and support the product(s) within target markets and around the world. Aspects such as vertical-market sales teams and specific vertical-market solutions are also considered.

**Market Responsiveness and Track Record:** This criterion evaluates vendors in light of the fact that market demands change over time. The availability of new products, services or licensing approaches in response to recent demands and the ability to recognize meaningful trends early in the adoption cycle are particularly important. A diversity of delivery models, as demanded by the market, is also important — for example, the ability to offer dbPaaS, software solutions, data warehouse "as a service" offerings and certified configurations.

**Marketing Execution:** This criterion evaluates the ability to generate and develop leads, channel development through Internet-enabled trial software delivery, and secure partnering agreements (including co-seller, co-marketing and co-lead management arrangements). Also considered are a vendor's participation in competitive situations, ability to expand its DMSA footprint within its customer base, and ability to win new customers. This criterion also considers the ratio of license sales to support and maintenance revenue as a measure of footprint expansion.

**Customer Experience:** This criterion evaluates vendors on the basis of a survey of reference customers and discussions with users of Gartner's inquiry service during the six quarters to 31 December 2017. Also considered are a vendor's track record on proofs of concept, customers' perceptions of products, and customers' loyalty (this reflects their tolerance of a vendor's practices and can indicate their level of satisfaction). This criterion is sensitive to year-over-year fluctuations, based on customer experience surveys. Additionally, customer input regarding the application of products to limited use cases can be significant, depending on the success or failure of the vendor's approach in this market.

**Operations:** This criterion evaluates the alignment of a vendor's organization, as well as whether and how this enhances its ability to deliver. This criterion considers a vendor's ability to support clients throughout the world, around the clock and in many languages. Anticipation of regional and global economic conditions is also assessed.

**Table 1.** Ability to Execute Evaluation Criteria

<b>Evaluation Criteria</b>	
<b>Product or Service</b>	
Weighting	High
<b>Overall Viability</b>	
Weighting	Low
<b>Sales Execution/Pricing</b>	
Weighting	High
<b>Market Responsiveness/Record</b>	
Weighting	Medium
<b>Marketing Execution</b>	
Weighting	Medium
<b>Customer Experience</b>	

Weighting	High
<b>Operations</b>	
Weighting	Low

Source: Gartner (February 2018)

## Completeness of Vision

The Completeness of Vision criteria assess a vendor's ability to understand the functional capabilities needed to support DMSA environments, to develop a product strategy that meets the market's requirements, to comprehend overall market trends, and to influence or lead the market when necessary. A visionary leadership role is necessary for long-term viability of both products and companies. A vendor's vision may be demonstrated — and improved — by a willingness to extend its influence throughout the market by working with independent third-party application software vendors that deliver additional functionality for its DMSA environment. A successful vendor will be able not only to understand the competitive landscape for DMSAs, but also to shape its future.

**Market Understanding:** This criterion covers a vendor's ability to understand the market and to shape its growth and direction. It examines a vendor's core competencies in this market. It also considers awareness of new trends, such as the increased demand from end users for mixed data management and access strategies that match the growing variety of skills and roles; the changing concept of the data warehouse and analytics data management, including strategies for metadata management and distributed data management; and a vendor's position with regard to emerging technologies such as data lakes and multimodel DBMSs. In addition, understanding the different audiences for various categories of data and associated SLAs (compromise, contender and candidate; see Note 4) is crucial. Also essential is a demonstrable track record of altering strategy and tactical delivery in response to opportunistic segments and broader market trends.

**Marketing Strategy:** This criterion considers 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. It covers a vendor's responses to the market trends identified above and any offers of alternative solutions in its marketing materials and plans. Additionally, an understanding of current product focus, and how it may need to change to address future market requirements, is deemed critical.

**Sales Strategy:** This criterion encompasses all plans to develop or expand channels and partnerships that assist with selling. Sales strategy is especially important for young organizations as it can enable them to greatly increase their market presence, while maintaining lower sales costs (for example, through co-selling or joint advertising). This criterion also covers a vendor's ability to communicate its vision to its field organization and, therefore, to existing and prospective customers. Of particular interest this year are pricing innovations and strategies, such as new licensing options and flexibility — especially in support of cloud, on-premises and hybrid deployments — and the availability of freeware and trial software.

**Offering (Product) Strategy:** This criterion is clearly distinguished from product execution. It evaluates the roadmap for enhancing capabilities across all four primary use cases (see Note 1). It also includes expected functionality and assesses timetables for meeting new market demands as manifest in roadmaps and development plans for:

Support for a varied level of data latency, with a growing focus on streaming and continuous data ingestion and real-time access for analysis.

Support for multiple data types, including a semantic design tier and metadata management capabilities.

System and solution auditing and health management to ensure use case SLA compliance.

Static and dynamic cost-based optimization, with the potential to span processing environments, data structures and storage options.

Management and orchestration of multiple processing engines.

Elastic workload management and process distribution across cloud, cloud and premises. This includes separation of processing and storage.

Support for a "best fit" engineering approach, for which have to support an open approach (enabling easy combination of technologies from different vendors), as an alternative to an integrated approach (demonstrating value from vendor stack integration) – or both approaches in combination.

**Business Model:** This criterion evaluates how a vendor's model of a target market suits its products and pricing, and whether the vendor can generate profits with this model – judging by its packaging and offerings. We consider reviews of publicly announced earnings and forward-looking statements relating to an intended market focus. For private companies, and to augment publicly available information, we use proxies for earnings and new-customer growth – such as the number of Gartner clients indicating interest in, or awareness of, a vendor's products during calls to our inquiry service.

**Vertical/Industry Strategy:** This criterion assesses a vendor's ability to understand its clients. A measurable level of influence within end-user communities and certification by industry standards bodies are important. A product or solution roadmap to support a specific industry is advantageous.

**Innovation:** This criterion assesses a vendor's development of new functionality, allocation of R&D spending, and leading of the market in new directions. Also addressed is the maturation of alternative delivery methods, such as cloud infrastructures, as well as solutions for hybrid on-premises-and-cloud and cloud-to-cloud data management support. A vendor's awareness of new methodologies and delivery trends, such as flexible pricing models and distributed data management, is also considered. Organizations are increasingly demanding data storage strategies that balance cost with performance optimization, so solutions that offer separation of compute and storage, or that address the age and "temperature" of data, will become increasingly important.

**Geographic Strategy:** This criterion considers a vendor's ability to address customer demands in different regions of the world using direct/internal resources alone or in combination with those of subsidiaries and partners. It also evaluates a vendor's global reach and its roadmap for addressing regulatory requirements in specific countries, particularly for cloud deployments.



**Table 2.** Completeness of Vision Evaluation Criteria

<b>Evaluation Criteria</b>	
<b>Market Understanding</b>	
Weighting	High
<b>Marketing Strategy</b>	
Weighting	Medium
<b>Sales Strategy</b>	
Weighting	Medium
<b>Offering (Product) Strategy</b>	
Weighting	High
<b>Business Model</b>	
Weighting	Low
<b>Vertical/Industry Strategy</b>	
Weighting	Low
<b>Innovation</b>	
Weighting	High
<b>Geographic Strategy</b>	
Weighting	Medium

Source: Gartner (February 2018)

## Quadrant Descriptions

### Leaders

The Leaders quadrant includes five traditional large vendors that have historically dominated this market, plus AWS, a market disruptor that entered the Leaders quadrant last year and has since solidified its position. Each of the five traditional vendors has had to adapt to changing market

conditions as the scope and breadth of DMSA platforms has shifted beyond the traditional data warehouse use case. The three leaders in terms of Ability to Execute — Oracle, AWS and Microsoft — have all invested heavily in the cloud, which has become increasingly important as a means of delivery. The other vendors in this quadrant have struggled to adapt to the new market conditions, and their execution has declined as a result, although their expertise and capabilities in addressing large complex deployments provide a solid foundation.

## Challengers

The two Challengers are smaller vendors that are executing well in their targeted customer segments. Snowflake has grown strongly, thanks to its cloud-based data-warehouse-as-a-service offering, and it plans to expand its cloud footprint to other providers, most notably Microsoft (Azure). MemSQL has embraced the concept of low-latency, real-time data warehousing through a single platform that can address multiple use cases. Both have focused on their execution in the past year, and this is reflected in their customer growth, revenue growth and increasing product maturity.

## Visionaries

Google has joined MarkLogic in the Visionaries quadrant, although these vendors have different visions for the market. MarkLogic continues to focus on frictionless data integration's potential to serve as the semantic reconciliation tier between data sources and various processing engines. Its vision has met with greater acceptance in the market as the focus on managing distributed data management platforms in an LDW context has increased. Google differentiates itself from its primary cloud competitors, AWS and Microsoft (Azure), through its focus on productizing technologies that have been used internally at Google for years. Unencumbered by "legacy" approaches, its focus remains serverless and consumption-based usage models for cloud-based DMSAs. The market has yet to fully align with the Visionaries — a situation that contributes to their comparatively low positions for Ability to Execute.

## Niche Players

The crowded Niche Players quadrant includes vendors that are battling to become Visionaries, Challengers and Leaders for the first time. It also includes vendors that have already made forays into adjacent quadrants, only to rebound to the Niche Players quadrant when market conditions forced a shift in focus. The Hadoop-based vendors, the Visionaries of 2016, fall into the latter group, as the hype around Hadoop continues to decline, forcing them to focus more on use cases for which their technology is a good fit (see "Hype Cycle for Data Management, 2017" ). Neo4j with its focus on new graph processing models, Qubole with its fresh approaches to Hadoop management, and reentrants to the market, like Actian, are in this group. Emerging China-based vendors complete the Niche Players: Alibaba Cloud, GBase and Huawei, the last two with platforms that combine an MPP relational database and Hadoop in LDW offerings. All these China-based vendors have had early success in their home market, but progression beyond the Niche Players quadrant will depend heavily on their ability to penetrate overseas markets, while demonstrating an ability to support a breadth of DMSA use cases.

## Context

Examining the placement of vendors in this year's Magic Quadrant, we see a pronounced gap between vendors on the left side bounded (to the right) by Snowflake and Pivotal, and those on the right bounded (to the left) by Google. The megavendors in the Leaders quadrant have

improved their vision and begun to pull away from the rest, along with the Visionaries Google and MarkLogic. The left side of the Magic Quadrant is populated by vendors whose vision is either undifferentiated in significant ways or insufficiently embraced by the market. Most vendors in the left-hand group are either seeking a unique vision to propel them into the Visionaries quadrant or focusing on executing particularly well in their niche markets, which may raise them to the status of Challengers.

Even within the Leaders quadrant, there is potential for vendors to spread out as they focus on either execution at the expense of vision or vice versa, which would push these vendors toward the Challengers of Visionaries quadrants. This possibility raises the potential for real innovation to enter the Leaders quadrant as space opens up in the middle, assuming the Leaders move left or down.

Finally, the white space at the top of the Magic Quadrant shows unrealized execution potential. As data management platforms become increasingly distributed (see "Modern Data Management Requires a Balance Between Collecting Data and Connecting to Data" ), challenges arise in relation to effective and efficient access to data. A focus on metadata management capabilities, combined with data virtualization approaches and a cohesive view of the LDW, is critical to the long-term success of such platforms. Although most of the Leaders have proposed solutions to address these needs, these solutions tend to be vendor-specific and lack the broad applicability that will be required for widespread adoption.

## Market Overview

The DMSA market is increasingly polarized. On the one hand, there is tremendous hype about new data types, new technologies to store and manage them efficiently, and new roles and skills to use them effectively. On the other, there is a recognition that investment in foundational traditional technologies (employed for the traditional data warehouse use case) will be essential to serve as a platform for the next wave of innovation. Both the traditional and the new are required for a modern DMSA platform (see "Survey Analysis: New Data and New Analytics Are All Mythology Unless You Add Skills" ).

As a result, the modern DMSA platform is coalescing around a new definition of normal capabilities. This includes support for multiple data models and data types, in-memory DBMSs and extended capabilities like graph processing – in what amounts to an LDW approach.

Gartner sees several key trends in the market (for more details, see "Data Management Solutions for Analytics: Current and Future States, 2017" ).

### Rise of Distributed Data Management Environments

DMSA architectures are increasingly based on multiple data repositories:

**Data warehouses:** These contain data from multiple sources, which is fully curated and cleansed.

**Data lakes:** These are typically distinguished by the storage of data in its native format and by the use of flexible schemas, often implemented "on read" and typically with light governance (though requirements for governance and data quality should not be ignored).

Data lakes are appropriate for batch operations, such as ETL, and for use as "landing pads" for incoming data that does not yet have fully defined business value. They can also serve as laboratories for data discovery and data science.

The LDW continues to gain traction as the target architecture for advanced analytics. It has reached greater than 15% adoption in its target market, up nearly 50% year on year (see "Survey Analysis: New Data and New Analytics Are All Mythology Unless You Add Skills" ). We expect this trend to continue.

Behind the rise of LDW architecture lies the recognition that a single data persistence tier is generally insufficient for the increased data and analytics demands that most organizations are experiencing. As a result, analytics-savvy organizations are augmenting their traditional relational data warehouse environments with nonrelational technologies like Hadoop, SQL-accessible cloud-based object stores, flexible deployment models, data virtualization technologies, and separation of storage and compute resources. Additionally, new pricing models — mostly in the cloud but some on-premises — are lowering the entry costs for DMSA environments (see "New Pricing Models for Cloud DBMSs Provide Cost Optimization Opportunities" ).

### **Further Cloud and Hybrid Cloud Deployments**

The cloud has established itself as a Tier 1 platform for DMSA deployment. Pure-play cloud vendor AWS appears in the Leaders quadrant for a second year, and all the Leaders are investing heavily in the cloud. Snowflake has emerged from the Niche Players quadrant as a Challenger, thanks to strengthened execution. New entrants to the Magic Quadrant include Alibaba Cloud, Qubole and Treasure Data, each with pure-play cloud offerings.

Although adoption of pure-play cloud solutions is growing strongly, on-premises deployments still predominate. Vendors with a credible hybrid cloud story will have an advantage as customers augment their on-premises solutions with cloud capabilities (see "Predicts 2018: Data Management Strategies Continue to Shift Toward Distributed" ).

### **Rise of China-Based Vendors**

China's market for DMSAs is huge, and a number of Chinese vendors have addressed it with significant success. They are now looking to branch into international markets. China-based vendors in this year's Magic Quadrant include Alibaba Cloud, GBase and Huawei. Chinese vendors that did not meet this year's inclusion criterion requiring more than 10% of production customers to be outside their home market include Eastern Jin Technology (Seabox Data) and Transwarp Technology, both with interesting approaches to using a Hadoop core. Gartner will continue to follow the progress of these vendors.

### **Contracting Use Cases for Hadoop-Only DMSAs**

As the hype around Hadoop declines, use cases for this technology are becoming more focused. We believe that Hadoop-based offerings are being pushed to the edges of the market. They are selected mainly for their most suitable use cases (such as ETL, data exploration and data science) and as queryable archives for less frequently accessed data. Hadoop-only vendors' efforts to address traditional relational processing are not working for them.

## **Evidence**

Gartner uses multiple sources of information to establish the positions of vendors in Magic Quadrants. These sources are adjusted to account for the maturity and size of a given market, and other factors.

For this Magic Quadrant, we used the following sources:

Original Gartner-published research, often including market share forecasts to establish the size of a market.

Publicly available data, such as earnings statements, partnership announcements, product announcements and published customer cases.

Gartner client inquiry service data collected from more than 17,000 inquiry interactions involving the authors and the wider Gartner analyst community during the 20 months to 31 December 2017. These interactions with clients provided information about, for example, use cases, issues encountered, license and support pricing, and implementation plans.

RFI surveys issued to vendors, in which they were asked to provide details about versions, release dates, customer counts and distribution of customers worldwide, among other things. Vendors could refuse to provide information in this way, at their discretion.

A survey of vendors' reference customers for DMSAs. Vendors were asked to submit contact details of 10 to 15 reference customers that generally reflected the requirements of the inclusion criteria. We then invited each reference customer to complete a 25-minute online survey. A total of 595 reference customers from 25 vendors completed the survey on 16 October 2016. The survey included questions to confirm that the customers were current license holders. Additionally, customers were asked to provide information about:

The size and scope of their implementations (especially in the case of open-source utilization).

Issues and software bugs.

Their feelings about their experience of the vendor, their use of other software tools in the environment, the types of data involved, and the rate of data refresh or load.

Their deployment plans.

- Current-year survey responses were used for all the commentary in this Magic Quadrant. Prior survey responses were used to identify trends only.
- Data from vendors' reference customers does not constitute a knowledge base representative of the whole DMSA market. The 595 reference customers are not representative of all customers in the DMSA market. Rather, they are customers whom selected vendors identified to Gartner and who agreed to participate in the survey. (For demographic profiles of the reference customers who participated, see Note 5.)

Gartner customer engagements, in which we provided specific support, were aggregated and anonymized to add perspective to the other, more expansive research approaches.

It is important to note that this was qualitative research that formed a cumulative base on which to form the opinions expressed in this Magic Quadrant.

## Note 1

### Primary DMSA Use Cases

#### **Traditional Data Warehouse**

This use case involves managing structured historical data from multiple sources. Data is mainly loaded through bulk and batch loading.

This use case can manage large volumes of data and is primarily used for standard reporting and dashboarding. To a lesser extent, it is also used for free-form ad hoc querying and mining, and for operational queries. It requires strong capabilities for system availability and administration and management, given the mixed workload's requirements for queries and user skills.

### **Real-Time Data Warehouse**

This use case adds a real-time component to analytics use cases, with the aim of reducing latency – the time lag between when data is generated and when it can be analyzed. It primarily manages structured data that is loaded continuously via microbatching and/or streaming analytics in support of real-time decision support, embedded analytics in applications, real-time data warehousing and operational data stores.

This use case primarily supports reporting and automated queries for operational needs and low-latency decision support (as in the HTAP use case described in "Magic Quadrant for Operational Database Management Systems" and "Critical Capabilities for Operational Database Management Systems" ). It requires high-availability and disaster recovery capabilities to meet operational demands. Managing different types of user and workload, such as ad hoc querying and mining, and the ability to store large volumes of historical data are of less importance, as the major driver is provision of a low-latency real-time view and analytics on operational data.

### **Context-Independent Data Warehouse**

This use case allows exploration of new data values, variants of data form and new relationships. It supports search, graph and other advanced capabilities for discovering new information models.

This use case is used primarily for free-form queries to support forecasting, predictive modeling or other mining styles, as well as queries supporting multiple data types and sources. It has no operational requirements and favors advanced users, such as data scientists and business analysts. It results in free-form queries across potentially multiple data types.

### **Logical Data Warehouse**

This use case manages data variety and volume for both structured and other content data types where the DBMS acts as a logical tier to a variety of data sources.

Besides structured data from transactional applications, this use case includes other content data types, such as machine data, text documents, images and videos. Because additional content types can produce large data volumes, and have specific data persistence requirements, access to data in disparate repositories is important. The LDW also has to support diverse query capabilities and diverse user skills. This use case supports queries reaching into sources other than the data warehouse DBMS alone, and may include metadata or data virtualization components.

## **Note 2 Industry Sectors**

Accommodation and food services

Administrative and support and waste management and remediation services

Agriculture, forestry, fishing and hunting

Arts, entertainment and recreation  
Construction  
Educational services  
Finance and insurance  
Healthcare and social assistance  
Information  
Management of companies and enterprises  
Manufacturing  
Mining  
Professional, scientific and technical services  
Public administration  
Real-estate rental and leasing  
Retail trade  
Transportation and warehousing  
Utilities  
Wholesale trade

### Note 3 Geographic Regions

North America (Canada and the U.S.)  
Latin America (including Mexico)  
Europe (Western and Eastern Europe)  
Middle East and Africa (including North Africa)  
Asia/Pacific (including Japan)

### Note 4 Data Categories and Associated SLAs

**Compromise.** There is agreement that the data model should be persisted for pervasive use by many end users, and that it exhibits a general tolerance for latency (even as low as two minutes). This SLA has two primary objectives: optimized performance and end-user comprehension of the data model. It is a compromise because the model is deployed to satisfy a low common denominator of use cases and specifically ignores exceptions. This is the traditional data warehousing approach, and is generally best thought of as "least common denominator" data for commonly shared analytics.

**Contender.** There is no agreement that any combination of data will be persisted or widely applicable to use cases. The result is a sense of transient data combinations being used by a diverse set of users. However, the source of the information is generally agreed to be an

adequate representation for exploring how the data can be used. Because of the changing and evolving nature of this SLA, zero-latency is often requested — but this is actually a proxy for seeing the data in as close to its native form as possible. This is data federation and is often supported by data virtualization software and even by multiple data marts. This approach is used when analysts have not reached agreement about how to combine disparate data, but seek to combine it under multiple models.

**Candidate.** There is wide access to the data asset and proposed model, but, because the data structure is complex and not always consistent, the default is to present multiple schema-on-read scenarios for different types of analysis. These scenarios explore unexpected forms in the data, but also create postulates of multiple alternative forms in parallel analytics from the data. This is big data analytics. A different barrier to entry exists here, in that users must understand how to parse and process data — almost like a DBMS. Candidates are suggested forms of reading data and potential uses of how that data is read. As such, they are submitted for consideration, but are not even contenders yet.

## Note 5 Profiles of Reference Customers

Numbers of reference customers reporting the following primary industry classifications for their organization:

- Services: 139
- Banking: 120
- Manufacturing and natural resources and mining: 50
- Media: 44
- Communications: 40
- Retail: 39
- Government: 26
- Education: 16
- Utilities: 16
- Healthcare: 12
- Transportation: 12
- Other: 31

Number of reference customers identifying their company as being in the following revenue size categories (U.S. dollars):

- Less than \$50 million: 80
- \$50 million to \$250 million: 77
- Over \$250 million to \$500 million: 30
- Over \$500 million to \$1 billion: 37
- Over \$1 billion to \$3 billion: 64



Over \$3 billion to \$10 billion: 70

Over \$10 billion to \$30 billion: 48

Over \$30 billion: 79

Number of reference customers identifying their organization as being in the government sector, public sector or education sector in the following employee size categories:

Fewer than 5,000 employees: 17

5,000 to 50,000 employees: 30

More than 50,000 employees: 13

## Evaluation Criteria Definitions

### Ability to Execute

**Product/Service:** Core goods and services offered by the vendor for 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:** Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will 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.

**Market Responsiveness/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 initiatives, 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 to 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 products 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 sets 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 vertical markets.

**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.

© 2018 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. or its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. If you are authorized to access this publication, your use of it is subject to the Usage Guidelines for Gartner Services ([/technology/about/policies/usage\\_guidelines.jsp](/technology/about/policies/usage_guidelines.jsp)) posted on gartner.com. The information contained in this publication has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This publication consists of the opinions of Gartner's research organization and should not be construed as statements of fact. The opinions expressed ([https://www.gartner.com/technology/contact/become-a-client.jsp?cm\\_sp=bac--reprint--banner](https://www.gartner.com/technology/contact/become-a-client.jsp?cm_sp=bac--reprint--banner)) herein are subject to change without notice. Gartner provides information technology research and advisory services to a wide range of technology consumers, manufacturers and sellers, and may have client relationships with, and derive revenues from, companies discussed herein. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "Guiding Principles on Independence and Objectivity. ([/technology/about/ombudsman/omb\\_guide2.jsp](/technology/about/ombudsman/omb_guide2.jsp))"

About (<http://www.gartner.com/technology/about.jsp>)

Careers (<http://www.gartner.com/technology/careers/>)

[Newsroom \(http://www.gartner.com/newsroom/\)](http://www.gartner.com/newsroom/)

[Policies \(http://www.gartner.com/technology/about/policies/guidelines\\_ov.jsp\)](http://www.gartner.com/technology/about/policies/guidelines_ov.jsp)

[Privacy \(https://www.gartner.com/privacy\)](https://www.gartner.com/privacy)

[Site Index \(http://www.gartner.com/technology/site-index.jsp\)](http://www.gartner.com/technology/site-index.jsp)

[IT Glossary \(http://www.gartner.com/it-glossary/\)](http://www.gartner.com/it-glossary/)

[Contact Gartner \(http://www.gartner.com/technology/contact/contact\\_gartner.jsp\)](http://www.gartner.com/technology/contact/contact_gartner.jsp)



[https://www.gartner.com/technology/contact/become-a-client.jsp?cm\\_sp=bac\\_-\\_reprint\\_-\\_banner](https://www.gartner.com/technology/contact/become-a-client.jsp?cm_sp=bac_-_reprint_-_banner)