G00297359

Critical Capabilities for Operational Database Management Systems

Published: 6 October 2016

Analyst(s): Merv Adrian, Donald Feinberg, Nick Heudecker, Terilyn Palanca, Rick Greenwald

New products for operational DBMS use cases are maturing, challenging data and analytics leaders to select the right technology. This research and companion Magic Quadrant report combine vendor capabilities and client implementation experience to identify products that meet requirements.

Key Findings

- Operational DBMS use cases include traditional transactions, distributed variable data, lightweight events and observations, and hybrid transactional/analytical processing (HTAP).
- Emerging multimodel, key-value, document, graph and wide-column DBMSs are challenging relational database management systems (RDBMSs), creating opportunities for best-fit engineering, and disrupting established company standards.
- Incumbent vendors are responding by adding multimodel capabilities, but hedging their bets (and confusing buyers) by also marketing specialized products.

Recommendations

Analytics leaders should:

- Classify the use cases under consideration and map them to the costs, deployment options and skills requirements of the products evaluated here.
- Avoid paying for features that are not needed right now as best-fit choices can be dramatically less costly if the necessary development and operational skills are available.
- Issue RFPs (using Gartner's RFP Toolkit templates) and conduct proof of concept (POC) exercises to confirm and validate candidates.
- Update standards, training and support organizations to manage, qualify and accommodate additional requirements for operational use cases if the company's technology portfolio is expanding.

Strategic Planning Assumptions

By 2017, all leading operational DBMSs will offer multiple data models, relational and nonrelational, in a single DBMS platform.

Through 2018, 50% of operational DBMS vendors with less than $50 million in annual revenue will disappear due to acquisitions, mergers or business failures.

By 2017, as "NoSQL" ceases to distinguish DBMSs, data and analytics leaders will select multimodel and/or specific document, key-value, graph and wide-column DBMSs.

What You Need to Know

Data and analytics leaders, architects and implementers can use this Critical Capabilities research as part of an evaluation of products based on what Gartner considers to be the nine most important functional capabilities across four major use cases (see Note 1). In tandem with the "Magic Quadrant for Operational Database Management Systems," which assesses the vendors who produce them, this facilitates choices that balance product and business relationship concerns. One operational DBMS product was chosen for vendors that have multiple offerings in their portfolios, guided by market leadership and representation across DBMS categories (see Note 2).

Our analysis synthesizes product information provided by vendors, information gathered from interactions with Gartner clients over the past 12 months, and relevant responses from our survey of the vendors' reference customers (852 responses), performed during May 2016.[1]

The ratings below, in which 3.0 represents "meets requirements," will enhance the breadth of information available to you, supporting a better-tuned decision process. Most scores achieve the meets requirements level. Nonetheless, any decision process you adopt should include POC tests with *your* data, on *your* equipment, and against *your* business requirements and SLAs.

Analysis

Critical Capabilities Use-Case Graphics

Ratings and summary scores range from 1.0 to 5.0, with 3.0 indicating meets requirements:

1 = Poor or Absent: Most or all defined requirements for a capability are not achieved

2 = Fair: Some requirements are not achieved

3 = Good: Meets requirements

4 = Excellent: Meets or exceeds some requirements

5 = Outstanding: Significantly exceeds requirements

Figure 1. Vendors' Product Scores for the Traditional Transactions Use Case

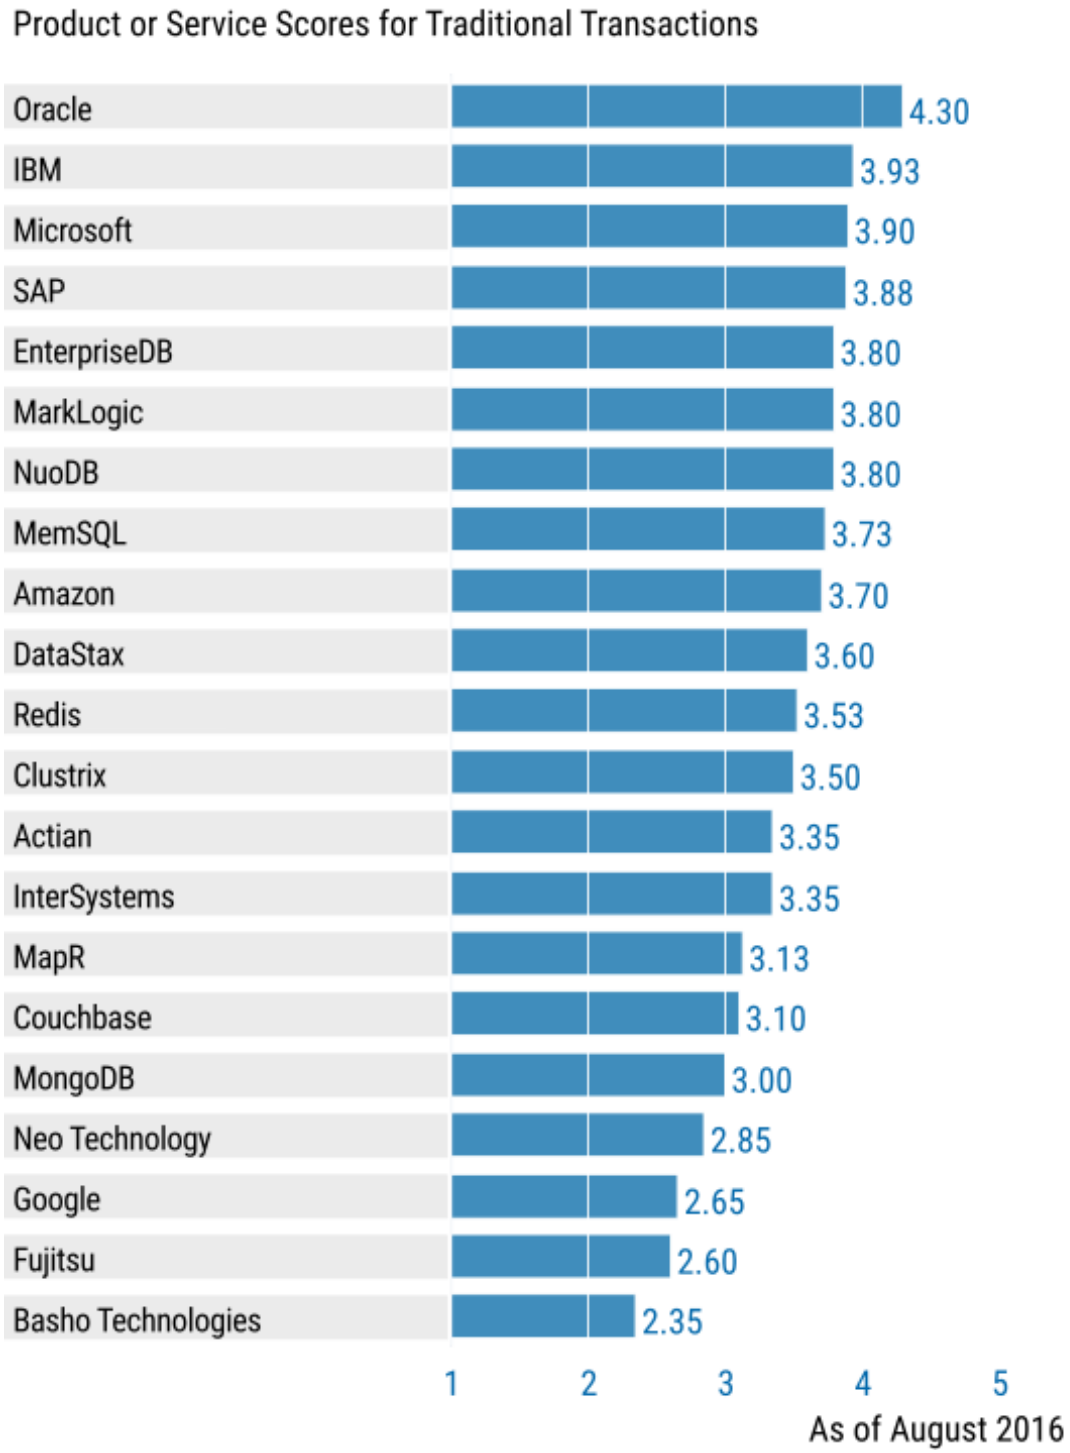

Source: Gartner (October 2016)

Figure 2. Vendors' Product Scores for the Distributed Variable Data Use Case

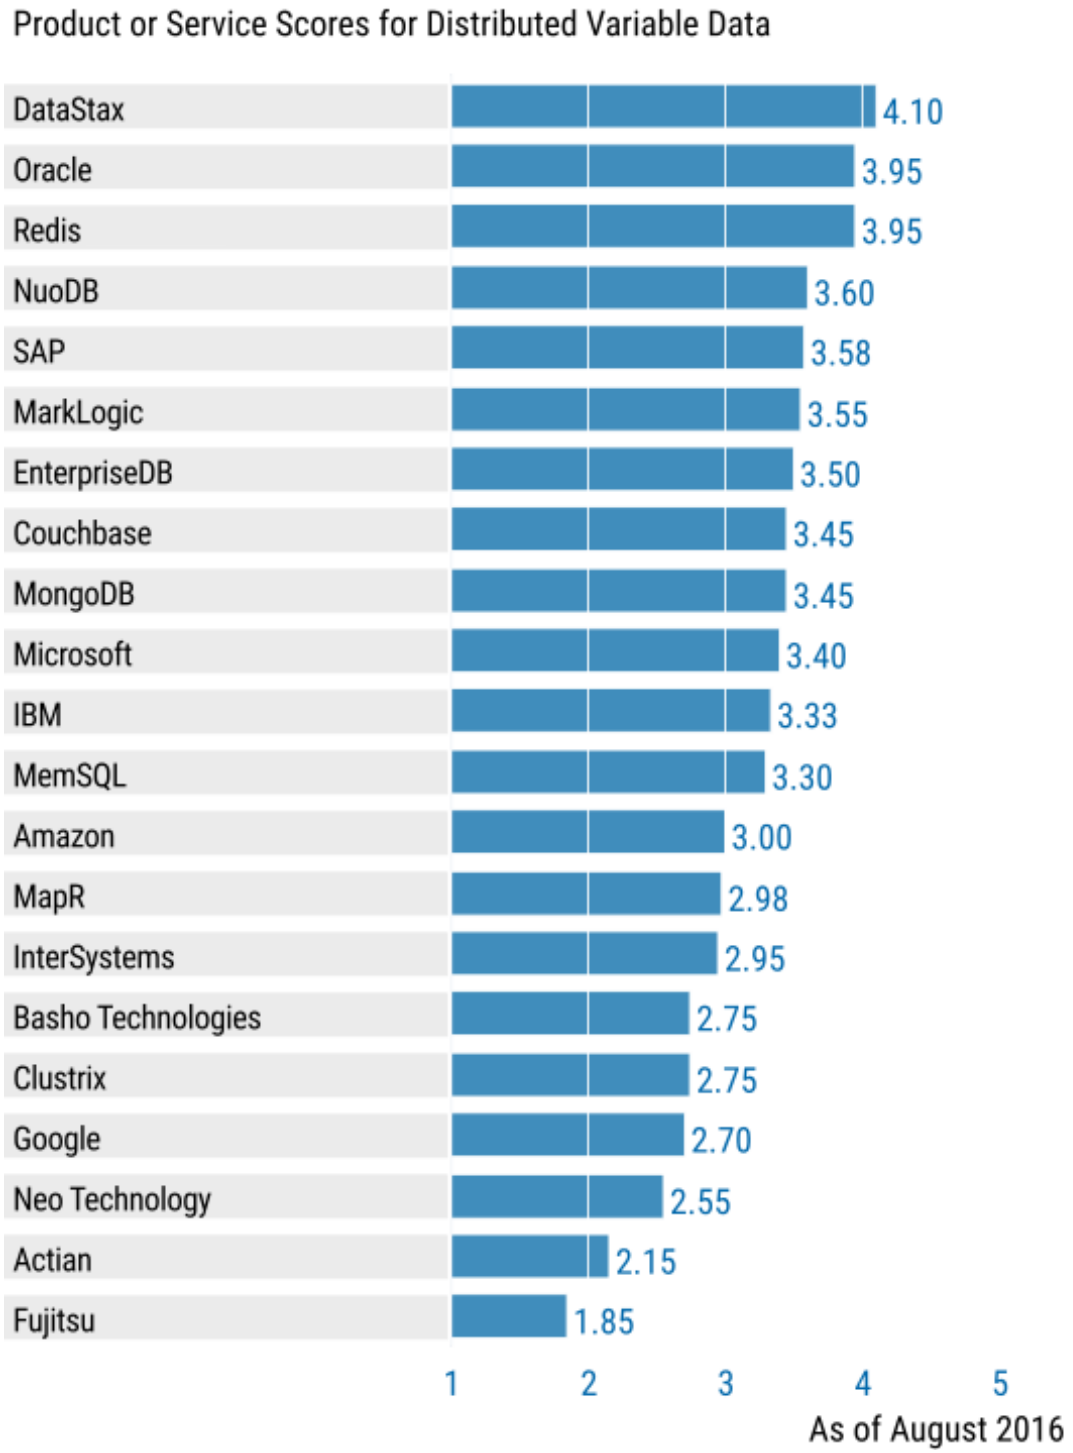

Source: Gartner (October 2016)

Figure 3. Vendors' Product Scores for Lightweight Events and Observations Use Case

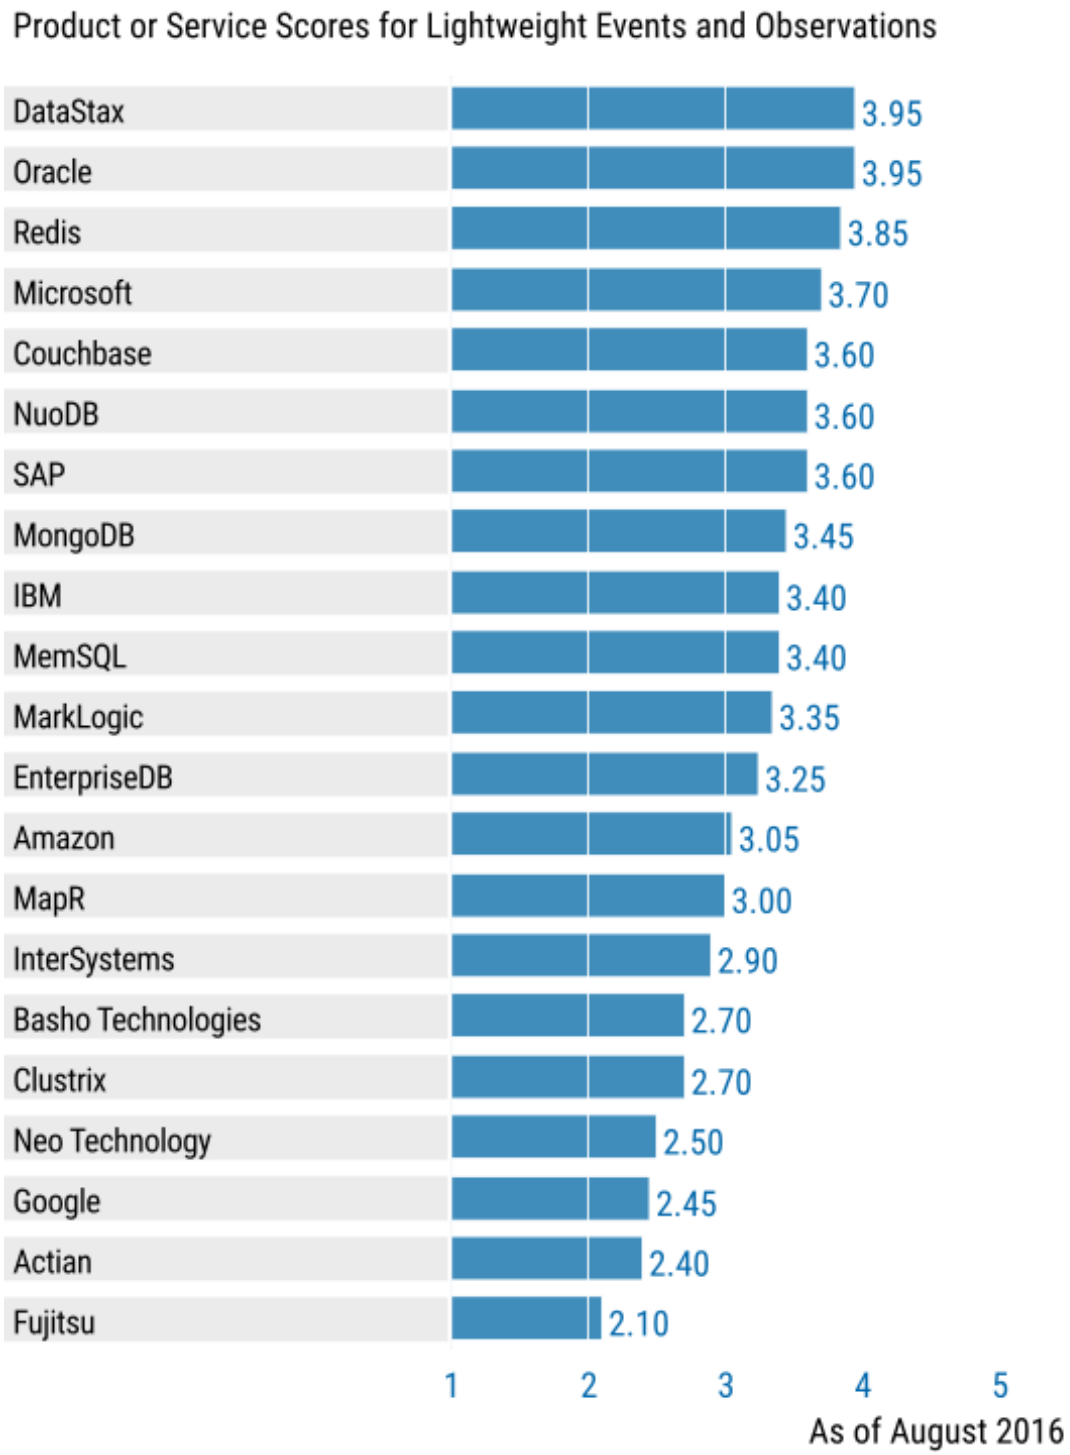

Source: Gartner (October 2016)

Figure 4. Vendors' Product Scores for Hybrid Transactional/Analytical Processing (HTAP) Use Case

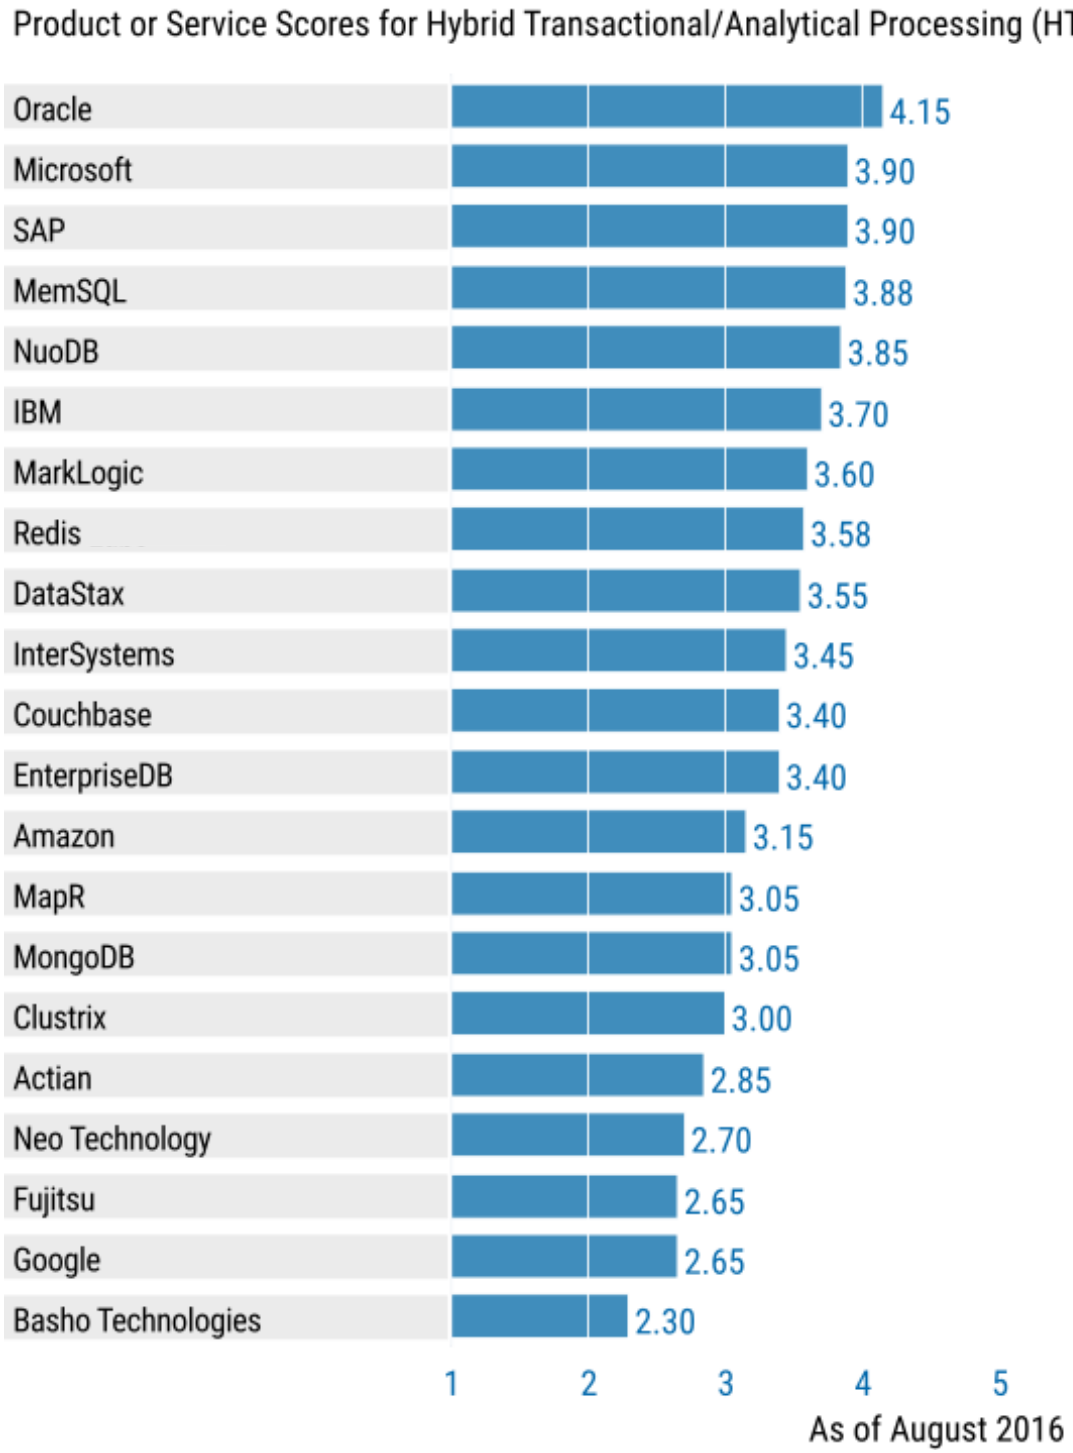

Source: Gartner (October 2016)

Vendors

Actian

Headquartered in Redwood City, California, Actian offers the closed-source Btrieve key-value DBMS, Ingres and PSQL RDBMS, and Versant object DBMS. Very few Ingres customers responded to Gartner's reference survey; on average, it had been in production 12 years in their shops. Reported database size and transaction rates were quite low, but overall performance rated above average. Actian Ingres meets requirements for traditional transaction use cases, but falls short for distributed variable data, HTAP and lightweight events and observations. Ingres significantly exceeds requirements for atomicity, consistency, isolation and durability (ACID) support. It received low ratings for tunable consistency and multimodel support; in the latter case Actian's portfolio is better described as polyglot — with a single interface for multiple separate engines.

Actian reports that only 12% of its customers use Ingres version 10.2, released in 2014, although surveyed reference customers rated the migration experiences "less complex than expected." Ingres received good survey scores on the frequency of problems encountered, although customers indicated difficulties integrating it with other DBMSs, despite the availability of Actian's DataCloud and DataConnect offerings. Customers reported low numbers of outages of two days or more, but shorter-term outages were significantly more frequent than the mean. However, high availability/ disaster recovery capabilities rated well. Support and professional services scores were low, and security was rated average (except for low scores for encryption).

Amazon

Amazon Web Services (AWS) is a wholly owned subsidiary of Amazon.com, headquartered in Seattle, Washington. AWS offers Amazon Aurora, as well as other DBMS services.

Closed-source, MySQL-compatible Amazon Aurora provides good ACID capabilities but suffers in its evaluation for tunable consistency. Amazon Aurora itself does not have much multimodel support, although AWS users would typically combine another AWS offering to meet this requirement.

This polyglot, best-fit approach limits the achievements of Amazon Aurora in many critical capabilities, since AWS views its set of DBMS services as best used together to provide different capabilities.

Basho Technologies

Headquartered in Bellevue, Washington, Basho Technologies offers Riak KV, a distributed, masterless key-value DBMS. Riak KV is available in open-source and commercial enterprise versions.

Riak KV scored low in the area of ACID support, which hurt it in the traditional transactions and HTAP use cases. For the distributed variable data and the lightweight events and observations use cases, this type of transaction support is not necessary, and the average to above-average scores

Riak KV received for high-speed ingest and processing, tunable consistency, automated data distribution and cloud/hybrid deployment capabilities raised its scores in these areas. Riak KV is a good example of a best-fit solution where that fit does not include scenarios where ACID transaction capabilities matter.

Clustrix

Headquartered in San Francisco, California, privately-held Clustrix offers a closed-source, low-administration, shared-nothing, distributed RDBMS, ClustrixDB, with replication and autoscaling. It is available as on-premises software and in the cloud.

ClustrixDB is a scale-out RDBMS compatible with MySQL, making it an attractive option if existing MySQL-backed applications are hitting scalability challenges. ClustrixDB is frequently used in e-commerce scenarios where horizontal scaling is important.

The DBMS offers ACID-compliant transactions across all nodes in the cluster. ClustrixDB scored well for traditional transactional use cases, placing roughly in the middle of the pack. The company recently added JSON and XML support, making it suitable for write-intensive lightweight events and observations as well as distributed variable data. ClustrixDB scored below average for HTAP capabilities. While it has ACID capabilities and scored well for high-speed ingestion and processing, it does not currently offer HTAP programming functionality beyond its streaming capabilities. Clustrix has in-memory capabilities on its roadmap for 8.0, but the nature of those capabilities is currently unknown.

Couchbase

Headquartered in Mountain View, California, privately held Couchbase offers Couchbase Server, an open-source multimodel DBMS featuring a distributed cache, a key-value DBMS and a document DBMS. Couchbase Server is available as a commercially supported enterprise edition, as well as an unsupported community edition. The company also offers Couchbase Mobile, a DBMS for mobile platforms with a synchronization gateway.

Couchbase Server's capabilities around cross-data-center replication and its multimodel features make it well suited for both lightweight events and observations, as well as distributed variable data. Typical use cases are multichannel profile management, personalization and content management, primarily in web and mobile environments. The product's characteristics also make it suitable for emerging Internet of Things (IoT) use cases. Couchbase Server had one of the highest reported transaction rates according to references.

Couchbase Server is also suitable for hybrid transaction/analytical processing use cases. The product has per-document ACID capabilities, making it suitable for some traditional transactional use cases, but it is not suited for use cases requiring cross-document ACID capabilities.

DataStax

Headquartered in Santa Clara, California, privately held DataStax provides DataStax Enterprise, a multimodel DBMS in an integrated platform built on the open-source Apache Cassandra DBMS

supporting tabular, key-value, document, and most recently a graph DBMS built with the Apache TinkerPop open-source project. The product is available in three editions: Basic, Standard and Max, which offer additional capabilities like integrated search, analytics, security, and management and monitoring tools. It is available on-premises, through various cloud providers, or as a hybrid cloud deployment.

DataStax Enterprise continues to be cited by customers for uses with distributed variable data needing high rates of ingest, as well as lightweight events and observations. This year's survey respondents cited transactional and HTAP uses more frequently, despite the product's middle-of-the-pack scores on ACID transactions and programmability for HTAP.

Although DataStax Enterprise can support limited ACID transactions through pinning, its limitations there, as well as in its support for the full SQL dialect in prevalent use today, make this product less suitable for traditional transactional environments and applications.

EnterpriseDB

Headquartered in Boston, Massachusetts, privately held EnterpriseDB ships EDB Postgres in Developer, Standard and Enterprise subscriptions, based on the PostgreSQL open-source DBMS. It is an operational DBMS standard in half of the surveyed customer shops, where it has typically been in production for four years. Half of surveyed customers are using release 9.5 (January 2016); the company reports 14% of all customers were doing so as of April 2016. EDB Postgres Advanced Server meets the requirements for all of our use cases, rating in the top five for both traditional transactions and distributed variable data. Its ACID support and multimodel capabilities were high points.

Most customers report low-end database sizes, transaction volumes and speeds, and rate its performance slightly above the mean. Half reported use of high availability/disaster recovery, high-speed ingest, and automated data distribution capabilities, and rated all highly. Last year's report cited steadily improving security and this year's scores reflected that, as EDB Postgres rated slightly above the mean in all security categories.

Low scores for programming features for HTAP cases will be addressed in upcoming releases with added support for in-memory tables and indexes, and foreign data wrappers (FDWs), supplementing its recent additions of JSON support and Integration Suite for Hadoop, MongoDB, and MySQL.

Fujitsu

Headquartered in Tokyo, Japan, publicly traded Fujitsu offers the Fujitsu Enterprise Postgres DBMS, based on the PostgreSQL open-source DBMS, along with an appliance (Fujitsu Integrated System Primeflex), a cloud platform (Fujitsu Cloud IaaS Trusted Public S5), the Symfoware Analytics Server, and the new MetaArc digital business platform. Fujitsu Enterprise Postgres is available on-premises or in the cloud.

Fujitsu Enterprise Postgres 9.5 meets the requirements for our traditional transaction and HTAP use cases, but falls short for distributed variable data and for lightweight events and observations. Its lowest ratings for high-speed ingest and processing and automated data distribution are key challenges, but not the only ones — Fujitsu scored below the midpoint on all criteria except cloud/hybrid deployment, ACID support, and programmability for HTAP, the latter significantly enhanced by the addition of an in-memory column store in 9.5.

Few Fujitsu customers responded to Gartner's reference survey. Their responses were very consistent in terms of use case — traditional transactions, with one HTAP use — but also in low satisfaction scores on product performance and ease of use among other factors. Database sizes and total number of transactions per day cited were lower than its competitors as well.

Google

Headquartered in Mountain View, California, Google is a wholly owned subsidiary of the publicly traded Alphabet holding company. Google offers closed-source Cloud Datastore, a managed schema-less document DBMS. This is the first year that Google Cloud Datastore is being evaluated for this Critical Capabilities research.

Google Cloud Datastore received a low score for cloud/hybrid deployment, since it has no on-premises offering. Because this capability is weighted across all use cases, and since Cloud Datastore did not receive any scores higher than average, the offering ranked in the lower portion of all use cases.

IBM

Based in Armonk, New York, publicly traded IBM offers closed-source IBM Compose, DB2 for Linux, Unix and Windows; DB2 for z/OS; DB2 on Cloud, dashDB for Transactions; IBM Graph, Information Management System; Informix; and open-source Cloudant and IBM Open Platform (Hadoop). IBM does not disclose customer numbers or the currency of deployed releases; 80% of surveyed customers said they were using the latest releases. Deployment models include hardware bundling, appliances, and deployment in IBM's SmartCloud and in third-party clouds.

DB2 is a mature, general-purpose ACID-capable RDBMS. It has proven capabilities for high-speed transactions. Scale-out online transaction processing (OLTP) clustering leverages in-memory caching using remote direct memory access (RDMA) while columnar storage and other features provide analytics query performance that assists in HTAP use cases. DB2 for z/OS routes analytics to the IBM DB2 Analytics Accelerator for cost-effective HTAP; this capability is not available as yet for other OS platforms, but is available in the cloud.

Customers surveyed for this research did not rate DB2 highly for high-speed ingest or automated data distribution, impacting the other use cases somewhat. DB2 meets requirements for all the use case highlighted here, and is in the top five for traditional transactions and HTAP.

Multimodel support includes XML and JSON (in LOBs for z/OS) and SQL functionality for JSON; the Linux, Unix and Windows versions support triplets with a Resource Description Framework (RDF) graph DBMS. Tunable consistency is still lacking, impacting distributed variable data cases, and

IBM's focus for events and observations is on its Informix DBMS, not discussed here. In general, IBM's portfolio is multimodel; its products less so.

InterSystems

Headquartered in Cambridge, Massachusetts, InterSystems was founded in 1978 and remains private. It markets Caché, which was originally an object-oriented DBMS, but is now a multimodel data platform. Caché has a strong position in the healthcare sector.

Caché scored within one point of the leader in each use case, putting it in the middle for all use cases (top 10 for HTAP). The company's positioning in the middle is likely due to the attention paid by other products to multimodel and cloud.

Caché is a strong DBMS with ACID compliance as well as some tunable consistency, and rates high for HTAP capabilities. Its cloud model remains based solely on hosting rather than platform as a service (PaaS).

MapR

Headquartered in San Jose, California, privately held MapR provides the MapR Converged Data Platform, which includes MapR-DB, MapR's Hadoop distribution, MapR-FS and MapR Streams. MapR-DB is an operational multimodel DBMS compatible with Apache HBase. It is available on-premises and through various cloud providers.

MapR-DB is a distributed DBMS specialized for modern data types, such as JSON. It can also support rich content types, such as audio, video and images. The DBMS supports ANSI-compatible SQL access via Apache Drill, which is included in the overall MapR Converged Data Platform. However, the DBMS' ACID capabilities are limited to a single row or entity within the DBMS. This limits its applicability to traditional transactional use cases, but it is suitable for lightweight events and observations. MapR-DB's selective geographic replication features also make it a good fit for use cases with distributed variable data.

MapR-DB can be used for HTAP use cases using Apache Spark to provide in-transaction analytics.

MarkLogic

Headquartered in San Carlos, California, privately held MarkLogic offers an ACID-compliant, closed-source document DBMS in Essential Enterprise, Global Enterprise and Mobile editions, and a free, fully featured developer version. It can be deployed via leading cloud and virtualization platforms, including those of AWS, Azure and VMware.

MarkLogic achieved the highest score possible for both ACID support and multimodel support. Since ACID support is rated highly for the traditional transactions and HTAP use cases, and multimodel support is rated highly for the distributed variable data and lightweight events and observations use cases, its rating for HTAP programmability improved over last year's. MarkLogic placed in the top seven in three of our four of use cases (see Figures 1 through 4).

MemSQL

MemSQL, founded in 2011, is a privately held company headquartered in San Francisco, California. It provides the closed-source MemSQL distributed in-memory ACID-compliant RDBMS, which includes an in-memory row store. In addition, MemSQL offers a memory- and disk-based column store, as well as Streamliner, its integrated Spark solution for real-time data pipelining.

MemSQL offers a relational DBMS featuring full compatibility with the MySQL protocol. Most customers use it for traditional transaction workloads and use cases. With its in-memory capabilities and the addition of Streamliner in 2016, MemSQL meets requirements today for both traditional transaction and hybrid transaction/analytic use cases, receiving the highest possible scores for both ACID support and programmability for HTAP, although below the mean for security and administration and management.

MemSQL also meets our requirements for distributed variable data and lightweight events and observations, although these currently form the minority of use cases cited by reference customers.

Microsoft

Headquartered in Redmond, Washington, publicly traded Microsoft markets its closed-source SQL Server DBMS for the operational DBMS market.

SQL Server is rated in the top four for three use cases: traditional transactions, lightweight events and observations, and HTAP. More than 95% of our surveyed customers use SQL Server for transaction processing and more than 40% use it for HTAP. SQL Server is increasingly used for enterprisewide, mission-critical applications. With the in-memory capabilities introduced in SQL Server 2012 and 2014, coupled with strong analytics (integration of R), it excels at HTAP. Its strong cloud offerings also make it an attractive choice for the lightweight events and observations use case.

SQL Server continues to enhance support for automated data distribution, rising slightly this year for the distributed variable data use case. The vision of stronger integration and compatibility with cloud offerings will help the use of multimodel and automated data distribution, making them even stronger in the future.

MongoDB

Headquartered in New York City, New York and Palo Alto, California, privately held MongoDB offers MongoDB Enterprise Advanced, a partially closed-source edition and an open-source version of its document DBMS. It also offers Atlas, a cloud version of its product, alongside cloud and on-premises management tools.

MongoDB scored at the median or above in all but one capability, with higher ratings in tunable consistency, automated data distribution, cloud/hybrid deployment and administration and management. MongoDB fell below the median for its support of ACID transactions, which affected its standing in both the traditional transaction and HTAP use cases.

The absolute scores for MongoDB improved in 6 of 10 categories over the 2015 rankings, with significant improvement in support for multimodel and programmability for HTAP.

Neo Technology

Privately held Neo Technology, headquartered in San Mateo, California, offers Neo4j — an open-source, native-graph DBMS — in Community, Enterprise and Government editions. Forty percent of customers surveyed said they are using Version 3, the most recent release, which shipped in February 2016. Cloud and hybrid deployment were found in nearly half of Neo4j's surveyed customers, one of the highest rates among vendors discussed here.

In a significant shift from last year's report, half of the surveyed customers cite Neo4j as their operational DBMS standard. On average, it has been in production for two years. Not surprisingly given their relative newness, the size, transaction volume and speed of Neo4j deployments are near the bottom of reported operational DBMS implementations.

Neo4j's graph model remains a specialized use case in and of itself, used for specific situations where its graph architecture excels. As the value of graph architecture becomes more widely appreciated, Neo4j should continue to see growth. Neo4j is close to meeting requirements for the traditional transactions use case, driven by its ability to handle transactions with ACID support, but falls short for other operational use cases. Its customers rated it very low for high availability/disaster recovery, automated data distribution, and security features including permission management, encryption and data masking. Neo's development roadmap will tackle several of these issues in upcoming releases.

NuoDB

Headquartered in Cambridge, Massachusetts, privately held NuoDB provides a highly available, geographically distributed (designed to span multiple data centers) closed-source operational SQL DBMS that scales horizontally and elastically. NuoDB supports on-premises, cloud and hybrid cloud deployments, and is available in Community and Enterprise editions. NuoDB is also available through the Amazon Cloud Marketplace.

NuoDB scored well in all four use cases (in the top six). This is due to the company's support for ACID and tunable consistency with strong high availability/disaster recovery and flexible configurations. Also, NuoDB's capabilities for geoclustering, multimaster data distribution, hybrid cloud and multimodel remain strong. NuoDB is tied for fourth place in lightweight events and third in observations and distributed variable data, making it one of the only vendors supporting ACID in the top of these use cases.

Oracle

Headquartered in Redwood Shores, California, publicly traded Oracle offers Oracle Database. Among its features and options, Oracle Database provides multimodel capabilities, in-memory analytics and robust high availability/disaster recovery features. In addition to software, several of

Oracle's DBMSs are available in engineered systems (appliances) as well as in public and private cloud infrastructure offerings.

Oracle Database is a general-purpose, closed-source relational DBMS that has evolved over the years to support new data types and structures (beyond pure relational). It has proven capabilities for high-speed transactions, security and database administration. While Oracle's cloud support is recent, its investment and focus is paying dividends in the form of increasingly robust cloud capabilities. In addition to ACID, Oracle Database also supports some tunable consistency and durability controls within a single database.

Oracle Database was either the first or second-ranked DBMS for each use case.

Redis

Headquartered in Mountain View, California, privately held Redis offers Redis Enterprise Cluster, an in-memory multimodel DBMS based on the open-source Redis project. In addition to its original key-value DBMS, it offers models for document, graph, and wide-column DBMS. It also supports various clustered and master-slave configurations in an on-premises environment, and is available as Redis Cloud, a fully-managed database-as-a-service offering.

Redis Enterprise Cluster performs well for use cases around lightweight events and observations, as well as distributed variable data. It is commonly used in a range of web, mobile and social applications to support features like caching, queues, time series and notification mechanisms.

Its in-memory foundation makes Redis Enterprise Cluster suitable for HTAP. ACID compliance is limited to a single cluster, and SQL support requires the use of Spark, limiting its suitability somewhat for traditional transactional use cases. Database sizes reported by reference customers were quite small relative to many competitors, but transactions per day were competitive.

SAP

Headquartered in Walldorf, Germany, publicly traded SAP markets closed-source SAP Hana. It is available as an appliance, as software only (Tailored Datacenter Integration) and in the cloud, including the Hana Cloud Platform.

SAP Hana is built for HTAP as an in-memory relational DBMS with strong consistency and high-speed transaction processing. It also supports a wide variety of data types and platform options in the cloud. SAP has added Hana Vora for connectivity to Hadoop and Spark, allowing data from Hadoop to be used in HTAP. Its weakest capability is tunable consistency.

SAP Hana scored higher in all of our use cases this year, with higher scores in many of the critical capabilities. In addition to HTAP (where it tied for second), it is a good choice for the other use cases.

Context

This Critical Capabilities report focuses on key DBMS capabilities for operational use of data. Gartner shifted some definitions and weights in this year's research, added a security capability, and combined high-speed ingest with high-speed processing. You can change our weightings for yourself on the interactive version of this document to suit your needs.

Both traditional relational DBMS leaders and insurgent players are emphasizing multimodel — incorporating specialized engines (including key-value, graph and document DBMS architectures for "unstructured data.") Organizations are leveraging these engines and pervasive adoption of in-memory computing to incorporate analytics directly into transactional processes for HTAP as well as to enable deployment to other new use cases, such as global scalability for web applications and emerging IoT applications involving event processing (see "Match Use Cases and Capabilities for Operational DBMSs").

Our weighting for multimodel support (termed "multiple data types/structures" in last year's Critical Capabilities report) has been lowered (as multimodel is becoming the industry standard) for all but the distributed variable data use case, where it was raised.

We focused more aggressively on tunable consistency this year — setting its weight to zero for traditional transactions, for example — on the premise that ACID is the very definition of this use case. The combined high-speed ingest and high-speed processing was weighted somewhat less than the combination of last year's two capabilities, and most highly for traditional transactions and for HTAP. The addition of security — based on reference survey results — highlighted some glaring weaknesses for some vendors and strengths for others.

Use this research to qualify candidates as you consider new projects. You may find opportunities for improvement beyond your present suppliers.

Product/Service Class Definition

The critical capabilities identified here address the major needs identified in this research. Operational DBMSs provide key DBMS capabilities for processing transactions, interactions, events and observations. In addition to "traditional" transactional uses such as ERP and financial systems, they are increasingly involved in adding analytics directly into transactional streams for HTAP, and in deploying other new use cases, such as global scalability for web applications and emerging IoT applications involving event processing.

Critical Capabilities Definitions

As part of the Magic Quadrant research process, we sought the views of vendors' reference customers via an online survey conducted in May 2016 (see the Evidence section). Our description of the capabilities denotes which were influenced significantly by survey responses.

High-Speed Ingest and Processing

Ability to load large volumes of data for high speed processing, often from multiple endpoints and in different formats (based on survey responses).

High-speed, high-volume processing can play a role in all use cases, but in varying forms, as it may not be similarly configured for all. For example, for many lightweight events and observations uses, continuous high-speed ingest is required, but substantial processing is not (based in part on survey responses).

ACID Support

Refers to the DBMS engine guaranteeing the properties of "atomicity, consistency, isolation and durability" that ensure reliable, recoverable database transactions.

Operational DBMS products differ in the scope in which they can guarantee a consistent view of transactions. A rating of 5.0 indicates consistency across multiple tables/objects and writers. A rating of 3.0 indicates a consistent view can be maintained for a smaller scope (such as a single table, object or collection) and thus "meets requirements" for specific cases, though not all. A rating of 1.0 indicates that reliable consistency cannot be guaranteed.

Tunable Consistency

This capability lets developers make runtime decisions on a per-operation basis, and on the level of required data consistency relative to a cluster of distributed DBMS nodes.

While ACID requirements may still be considered the gold standard, enterprises tell Gartner that many of their new applications do not require ACID, and that using tunable consistency simplifies design and implementation. Traditional vendors often have parameters that may be carefully tweaked to "turn off" hard-won consistency (and many have used these switches over the years to achieve performance goals for benchmarks). Newer vendors who began their designs with the assumption of relaxed consistency may have an easier-to-implement model as a result.

Multimodel Support

This capability represents the use of multiple optimized engines to support data types in addition to traditional structured data (such as graph, document and object relational).

This capability goes beyond simply providing a field that "can contain anything" and leaving its usage to the programmers working with the DBMS. Tools for managing and handling the data, interfaces for providing effective query, and integration methods for using the data with other stored formats are required.

Automated Data Distribution

Permits configuration and execution of the rules and policies for how and to what extent the DBMS engine distributes data throughout a cluster of distributed nodes (based on survey responses).

Once the norm, "self-administered" sharding and distribution strategies are now seen as complex and challenging. Distributed variable data use cases — especially geographically distributed ones — are on the rise among enterprises that view automated data distribution as a requirement, and who turn to vendors that make it a product feature, not a complex skill to be acquired.

Cloud/Hybrid Deployment

The ability to deploy the DBMS across on-premises and cloud environments synergistically.

This could include single applications spanning and synchronizing cloud and on-premises environments, or different use cases being deployed to different environments (development and testing in the cloud, or backup to the cloud, for example). Applications can move easily from one environment to the other.

Programmability for HTAP

Supports embedded analytics in applications as part of the processing of a transaction while sustaining high performance to meet SLAs, normally using a single database instance.

Leveraging in-memory computing for both analytics and transactions, and exploiting the multiple engines available in a multimodel architecture are key elements (see Note 3 for a definition of HTAP).

Security

Includes configuring and enforcing policy-based access (by role and attribute) to data, data masking and encryption policies (based on survey responses).

Administration and Management

Describes the product's tools and features to support implementation and operations, including upgrade and availability features. It is largely measured by customer reporting.

Scoring is also affected by the complexity of deployment and by vendor history. Some vendors have recent offerings for which upgrades may not have been released. Administration capabilities (such as role-based activities, advisors, utilization and capacity planning), resource allocation features, and user interface — as well as complexity of deployment and management — are all considered.

Use Cases

Traditional Transactions

Manages centralized transactions, with high-speed, high-volume data ingest and processing, ACID properties, and security weighted highly.

Typical examples include retail point-of-sale systems, financial trading platforms, and telecom service providers collecting structured data at high speed from the network and passing to billing systems for multiple carriers at high speeds. Fixed, stable schemas, high availability, strong in-memory-based read and write performance, and low latency are characteristic requirements.

Distributed Variable Data

Requires automated data distribution and management, multimodel support for several data types and structures, and flexible schemas that may change often.

Customer "360-degree view" examples are designed to source data from multiple legacy systems, with document-based details and transactions across different lines of business. Document DBMSs have been used by insurance companies to handle tens of millions of agreements and hundreds of millions of transactions at call centers across North America. The huge and lengthy effort required to normalize data into a relational format would result in a relatively inflexible and difficult-to-change system.

Lightweight Events and Observations

Captures IoT-style edge events, usually with a small local footprint, and cloud/hybrid deployment. May include staged processing and transmission to other process components.

Utility examples include wireless network metering that automatically collects and analyzes usage data from billions of simple data points. High availability and a relatively self-maintaining and easy-to-deploy infrastructure are typical requirements. Key-value DBMSs fit these requirements well. Relational DBMSs typically do not cope well with the quickly rising levels of revenue-critical data that are beyond the scope of simple file-based capture (see "DBMS Characteristics for the Internet of Things").

Hybrid Transactional/Analytical Processing (HTAP)

Enables analytics to affect processing through multiple states within transaction scope in a single database, while maintaining low latency and strong security.

Manufacturing company purchase order systems use analytics on inventory, estimating lead time for additional product, and factoring user buying history to offer discounted pricing in real time. Similarly, airline pricing of seats based on the current sales for that flight, and the desired yield for the flight, is enabled by a system that can base decisions on analytics, including the use of interaction data and predictive models of buyer behavior. Scoring of a customer's credit before completion, or variable manufacturing price quotes based on real-time material availability and cost are also examples. See Note 3 for a definition of HTAP.

Inclusion Criteria

In this Critical Capabilities report, we include one product from each vendor included in the "Magic Quadrant for Operational Database Management Systems."

Market presence: Vendor operational DBMS products must have production presence in a minimum of three industry sectors.

Industry sectors are:

- 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
- Transportation and warehousing
- Utilities
- Wholesale

Additionally, the vendor must have market presence in a minimum of two of the following geographic regions:

- North America
- Latin America
- Europe
- Middle East and Africa
- Asia/Pacific

Regional market presence is represented by dedicated sales offices or distribution partnerships in a specific region.

Software availability: Vendors must have operational DBMS software that has been generally available for licensing or supported download for at least a year (as of midnight, U.S. Eastern Daylight Time on 1 July 2016).

Software releases: We use the most recent generally available release of the software to evaluate current technical capabilities. We do not consider beta, "early access," "technology preview," "ramp up" or other not generally available releases. For customer references and reference survey responses, all versions currently used in production are considered. When older versions are in use, we consider whether later releases may have addressed reported issues, but also the rate at which customers move to newer versions.

Feature availability: Product evaluations include technical capabilities, features and functionality present in the product or supported for download (as of midnight, U.S. Eastern Daylight Time on 1 July 2016). Capabilities, product features or functionality released after this date can be included at Gartner's discretion and in a manner Gartner deems appropriate for ensuring the quality of our research product on behalf of our nonvendor clients. We also consider how such later releases can reasonably impact the end-user experience.

Customers and revenue: Vendors must generate a minimum of $50 million in verifiable annual DBMS software revenue or a minimum $15 million in estimated revenue with 50% year-over-year growth or minimum venture capital/private equity funding of $50 million. For information about venture capital investments in the operational DBMS vendors discussed in the Magic Quadrant, we consulted CrunchBase. Revenue can be from licenses, support and/or maintenance. Gartner may include additional vendors based on undisclosed references in cases of known use for classified but unspecified use cases.

Support: The vendor must provide support for these operational DBMS product(s). For open-source DBMS versions, maintenance and support must be available from a vendor that owns (or has substantial control over) the source code, and offers it under an open-source license (such as the GPL or ASL). However, only the core DBMS engine must be under the open-source license to classify as an open-source DBMS.

Services: Vendors participating in the operational DBMS market must demonstrate their ability to deliver the necessary services to support transaction systems via the establishment and delivery of support processes, professional services and/or committed resources and budgets.

Product breadth of use: Vendors must demonstrate support of production operational DBMS customers in at least *three* of the major geographic regions (North America, Latin America, Europe, Middle East and Africa, and the Asia/Pacific region). Additionally, vendors must demonstrate production use in a minimum of three industry segments (see the "Market Presence" criteria for a list of industry segments).

Excluded products: Some vendors are explicitly excluded from this Critical Capabilities research. Products that "add a layer" to (and require or embed a complete or near-complete implementation

of) another commercially marketed product such as MySQL, HBase or PostgreSQL are not included. Highly specialized engines such as embedded-only, text-only or object-only databases (which may perform some transactions to target small subsets of operational use cases) are also excluded. Operational DBMS products with over 50% of their production deployments embedded are also excluded. Products covered by technologies listed in the Replacement Phase of "IT Market Clock for Database Management Systems, 2016" will be excluded. Finally, "streaming" engines, whose use cases are dominated by immediate event processing, and which are rarely (if ever) used for subsequent management of the data involved, are also excluded.

Product categories specifically excluded from this Critical Capabilities are:

- Data-warehouse-only DBMS products
- Data grid products
- Complex-event processing or streaming data engines

Prerelational and object DBMS products are excluded from this research because Gartner considers them to be at the replacement phase of their maturity (see "IT Market Clock for Database Management Systems, 2016").

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants (and their companion Critical Capabilities) as markets change. As a result of these adjustments, the mix of vendors 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

The following vendors have been added to this Critical Capabilities report:

- Google

Dropped

The following vendors have been dropped from this year's Critical Capabilities report:

- FairCom (did not submit revenue numbers as part of Gartner's formal request for information).
- MariaDB (did not meet revenue requirements for inclusion).
- Percona (did not meet revenue requirements for inclusion).
- TmaxSoft (did not meet revenue requirements for inclusion).
- VoltDB (did not meet revenue requirements for inclusion).

Other Vendors to Consider

A number of other vendors may be worthy of consideration, depending on requirements, although they did not meet the criteria for inclusion in this year's Magic Quadrant, and thus are not assessed in this Critical Capabilities report.

Aerospike: Headquartered in Mountain View, California, and founded in 2009, Aerospike markets a hybrid in-memory/flash NoSQL DBMS — a real-time data platform — for the operational DBMS market. It is available both as an open-source community version and an Enterprise edition.

Altibase: Headquartered in Seoul, South Korea and Fort Lee, New Jersey, Altibase offers Altibase HDB, an SQL operational DBMS capable of using in-memory, solid-state disk, traditional disk and hybrid storage. Altibase also offers the Altibase XDB, an in-memory-only DBMS. Its products are available on-premises, on a virtualized basis and through cloud providers.

Cloudera: Headquartered in Palo Alto, California, Cloudera offers Cloudera Enterprise, a commercial version of Apache Hadoop for which Apache HBase provides the operational DBMS capabilities. Cloudera Enterprise is available both on-premises and through various cloud providers.

FairCom: Founded in 1979 and headquartered in Columbia, Missouri, and privately owned. FairCom c-treeACE (Advanced Core Engine), one of the oldest NoSQL DBMSs, is a fully ACID, key-value DBMS with both NoSQL (Indexed Sequential Access Method [ISAM]) interfaces and SQL. It supports transactions with an embedded or stand-alone engine.

MariaDB: Headquartered in Espoo, Finland, MariaDB markets MariaDB, an open-source, in-memory-capable, multimodel relational DBMS that is fully compatible with Oracle MySQL; MariaDB MaxScale, a database proxy for scaling MySQL and MariaDB; and MariaDB Enterprise, a commercially supported bundle with enterprise-targeted add-on components. All are available on Linux (where MariaDB is the default DBMS in the Red Hat and SUSE distributions) and on Microsoft Windows.

McObject: Headquartered in Issaquah, Washington, McObject offers eXtremeDB version 6.0, a small-footprint relational in-memory DBMS with extended array and vector support. Since 2001, millions of copies of eXtremeDB have been deployed worldwide in embedded and real-time applications.

OrientDB: Based in London, offers an open-source, multimodel operational DBMS that is predominantly graph with rich security features. OrientDB is available on-premises and through various cloud providers. References score OrientDB above average for security, particularly permission management.

Percona: Headquartered in Durham, North Carolina, Percona delivers enterprise-class, open-source versions of MySQL via its Percona Server and Percona XtraDB Cluster products, and of MongoDB via Percona TokuMX and Percona Server for MongoDB. Percona provides support, consulting and managed services for MySQL, MariaDB, Amazon RDS, OpenStack and MongoDB, as well as its own solutions.

TmaxSoft: Headquartered in Bundang-gu in the Seoul Capital Area, South Korea, TmaxSoft (formerly TmaxData), provides Tibero, an SQL relational DBMS featuring various clustering options, integrated encryption and compatibility with other vendors' DBMS products. It is available on-premises, in the cloud and via an appliance.

VoltDB: Headquartered in Boston, Massachusetts, VoltDB markets an open-source, in-memory, row-store operational relational DBMS for shared-nothing clusters. It is available in several package formats for leading cloud and virtualization platforms.

Other Changes

The 2016 Critical Capabilities assesses Amazon Aurora. In 2015, the product assessed was DynamoDB.

Table 1. Weighting for Critical Capabilities in Use Cases

Critical Capabilities	Traditional Transactions	Distributed Variable Data	Lightweight Events and Observations	Hybrid Transactional/Analytical Processing (HTAP)
High-Speed Ingest and Processing	25%	15%	20%	20%
ACID Support	25%	0%	0%	15%
Tunable Consistency	0%	10%	10%	5%
Multimodel Support	5%	20%	10%	5%
Automated Data Distribution	5%	20%	10%	5%
Cloud/Hybrid Deployment	15%	10%	20%	5%
Programmability for HTAP	0%	5%	10%	20%
Administration and Management	5%	5%	5%	5%
Security	20%	15%	15%	20%
Total	100%	100%	100%	100%
				As of August 2016

Source: Gartner (October 2016)

Critical Capabilities Rating

Table 2. Product/Service Rating on Critical Capabilities

Critical Capabilities	Actian	Amazon	Basho Technologies	Clustrix	Couchbase	DataStax	EnterpriseDB	Fujitsu	Google	IBM	InterSystems	MapR	MarkLogic	MemSQL	Microsoft	MongoDB	Neo Technology	NuoDB	Oracle	Redis	SAP
High-Speed Ingest and Processing	3.0	3.0	3.0	3.0	5.0	4.0	3.0	1.0	3.0	3.5	3.0	3.5	3.0	4.0	3.0	3.0	2.0	4.0	4.0	5.0	3.5
ACID Support	5.0	5.0	2.0	5.0	2.0	3.0	5.0	5.0	3.0	5.0	5.0	3.0	5.0	5.0	5.0	2.0	5.0	5.0	5.0	2.5	5.0
Tunable Consistency	1.0	2.0	4.0	1.0	5.0	5.0	3.0	2.0	3.0	1.0	3.0	2.0	3.0	3.0	3.0	4.0	3.0	5.0	3.0	5.0	2.0
Multimodel Support	1.0	2.0	3.0	3.0	3.0	5.0	5.0	2.0	3.0	4.0	3.0	3.0	5.0	3.5	3.0	3.0	3.0	4.0	5.0	3.0	5.0
Automated Data Distribution	2.0	4.0	3.0	3.0	3.0	4.0	3.0	1.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0	4.0	3.0	3.0	3.0	5.0	3.0
Cloud/Hybrid Deployment	3.0	4.0	3.0	3.0	3.0	4.0	3.0	3.0	1.0	4.0	2.0	3.0	3.0	3.0	5.0	4.0	3.0	3.0	4.0	3.0	4.0
Programmability for HTAP	2.0	2.0	2.0	2.0	4.0	3.0	2.0	4.0	2.0	4.0	4.0	3.0	3.0	5.0	5.0	3.0	3.0	4.0	4.0	3.0	5.0
Administration and Management	3.0	4.0	3.0	3.0	4.0	4.0	3.0	2.0	2.0	5.0	3.0	3.0	3.0	2.0	5.0	4.0	3.0	3.0	5.0	4.0	3.0
Security	3.0	3.0	1.0	3.0	2.0	3.0	4.0	2.0	3.0	3.0	3.0	3.0	4.0	3.0	3.0	3.0	1.0	3.0	4.0	3.0	3.0
As of August 2016																					

Source: Gartner (October 2016)

This methodology requires analysts to identify the critical capabilities for a class of products/services. Each capability is then weighed in terms of its relative importance for specific product/service use cases.

Table 3 shows the product/service scores for each use case. The scores, which are generated by multiplying the use-case weightings by the product/service ratings, summarize how well the critical capabilities are met for each use case.

Table 3. Product Score in Use Cases

Use Cases	Actian	Amazon	Basho Technologies	Clustrix	Couchbase	DataStax	EnterpriseDB	Fujitsu	Google	IBM	InterSystems	MapR	MarkLogic	MemSQL	Microsoft	MongoDB	Neo Technology	NuoDB	Oracle	Redis	SAP
Traditional Transactions	3.35	3.70	2.35	3.50	3.10	3.60	3.80	2.60	2.65	3.93	3.35	3.13	3.80	3.73	3.90	3.00	2.85	3.80	4.30	3.53	3.88
Distributed Variable Data	2.15	3.00	2.75	2.75	3.45	4.10	3.50	1.85	2.70	3.33	2.95	2.98	3.55	3.30	3.40	3.45	2.55	3.60	3.95	3.95	3.58
Lightweight Events and Observations	2.40	3.05	2.70	2.70	3.60	3.95	3.25	2.10	2.45	3.40	2.90	3.00	3.35	3.40	3.70	3.45	2.50	3.60	3.95	3.85	3.60
Hybrid Transactional/ Analytical Process-	2.85	3.15	2.30	3.00	3.40	3.55	3.40	2.65	2.65	3.70	3.45	3.05	3.60	3.88	3.90	3.05	2.70	3.85	4.15	3.58	3.90

Use Cases	Actian	Amazon	Basho Technologies	Clustrix	Couchbase	DataStax	EnterpriseDB	Fujitsu	Google	IBM	InterSystems	MapR	MarkLogic	MemSQL	Microsoft	MongoDB	Neo Technology	NuoDB	Oracle	Redis	SAP
ing (HTAP)																					
																					As of August 2016

Source: Gartner (October 2016)

To determine an overall score for each product/service in the use cases, multiply the ratings in Table 2 by the weightings shown in Table 1.

Gartner Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"IT Market Clock for Database Management Systems, 2016"

"How to Enable Digital Business Innovation via Hybrid Transaction/Analytical Processing"

"Are DBMS Appliances in Your Future? Don't Bet on It!"

"How Markets and Vendors Are Evaluated in Gartner Magic Quadrants"

"The OLTP DBMS Market Becomes the Operational DBMS Market"

"Match Use Cases and Capabilities for Operational DBMSs"

"Toolkit: TCO Calculator for Database Management Systems"

"The Rise of Polyglot Persistence Demands Your Consideration"

"The State of Open-Source RDBMSs, 2015"

"Market Share: All Software Markets, Worldwide, 2015"

"Market Guide for In-Memory DBMS"

"Market Guide for NoSQL DBMSs"

"Market Guide for Database Platform as a Service"

Evidence

[1] **Reference Customer Survey**
In addition to hundreds of interactions with users of Gartner's client inquiry service, as part of the Magic Quadrant and Critical Capability process, we sought the views of vendors' reference customers via an online survey. The survey included requests for feedback about vendor maturity (for example, understanding of industries, provision of innovation, responsiveness to new requests, TCO and pricing) and product capabilities (for example, flexibility in data modeling, support for data quality, UI support for data stewardship, internal workflow and support for multiple architectural styles). Over 500 organizations, representing all the featured vendors' reference bases, responded to the survey, which was held in May 2016. The reference customers were generally pleased with their vendors and products, but gave relatively low marks in some areas, which we have detailed in the analysis of each vendor. Some of the issues may be historical, because not all organizations are on the latest product versions.

Gartner's Client Inquiry Service Data

Gartner maintains an extensive database of information about all inquiries to our client inquiry service. Our information management team received over 4,600 inquiries during the Magic Quadrant research period of July 2015 to July 2016, of which over 900 were specifically about DBMSs. We used the sentiments apparent from these inquiries to formulate the opinions expressed in this Critical Capabilities.

CrunchBase was used for venture capital investment information.

Note 1 Definition of an Operational DBMS Workload

For the purposes of this evaluation, the workloads we expect to be managed by an operational DBMS include: Batch/bulk loading; real-time or continuous data loading; concurrent online and web-based new and update transactions; operational reporting; and management of externally distributed processes such as "look-aside" queries. Operational DBMS products must provide the ability to prioritize these multiple workloads to ensure SLAs are met when they operate concurrently.

Note 2 Definition of a Database Management System (DBMS)

Gartner defines a DBMS as a complete software system that supports and manages a database (or databases) in some form of storage medium (which can include hard-disk drives, flash memory, solid-state drives or even DRAM).

Note 3 Definition of HTAP

Hybrid transactional/analytical processing (HTAP) — at times referred to as hybrid online analytical processing (OLAP)/online transaction processing (OLTP), or hybrid OLTP and analytics — is an emerging application architecture that "breaks the wall" between transaction processing and analytics. It enables more informed and "in business real time" decision making. HTAP is defined as "an application architecture whereby concurrent analytical and transaction processing algorithms share the same data (and data infrastructure)." Therefore, HTAP allows advanced analytics to be run in real time on "inflight" transaction data, providing an architecture that empowers users to respond more effectively to business moments.

Critical Capabilities Methodology

This methodology requires analysts to identify the critical capabilities for a class of products or services. Each capability is then weighted in terms of its relative importance for specific product or service use cases. Next, products/services are rated in terms of how well they achieve each of the critical capabilities. A score that summarizes how well they meet the critical capabilities for each use case is then calculated for each product/service.

"Critical capabilities" are attributes that differentiate products/services in a class in terms of their quality and performance. Gartner recommends that users consider the

set of critical capabilities as some of the most important criteria for acquisition decisions.

In defining the product/service category for evaluation, the analyst first identifies the leading uses for the products/services in this market. What needs are end-users looking to fulfill, when considering products/services in this market? Use cases should match common client deployment scenarios. These distinct client scenarios define the Use Cases.

The analyst then identifies the critical capabilities. These capabilities are generalized groups of features commonly required by this class of products/services. Each capability is assigned a level of importance in fulfilling that particular need; some sets of features are more important than others, depending on the use case being evaluated.

Each vendor's product or service is evaluated in terms of how well it delivers each capability, on a five-point scale. These ratings are displayed side-by-side for all vendors, allowing easy comparisons between the different sets of features.

Ratings and summary scores range from 1.0 to 5.0:

1 = Poor or Absent: most or all defined requirements for a capability are not achieved

2 = Fair: some requirements are not achieved

3 = Good: meets requirements

4 = Excellent: meets or exceeds some requirements

5 = Outstanding: significantly exceeds requirements

To determine an overall score for each product in the use cases, the product ratings are multiplied by the weightings to come up with the product score in use cases.

The critical capabilities Gartner has selected do not represent all capabilities for any product; therefore, may not represent those most important for a specific use situation or business objective. Clients should use a critical capabilities analysis as one of several sources of input about a product before making a product/service decision.

GARTNER HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters
AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit http://www.gartner.com/technology/about.jsp