FORRESTER®

The Forrester Wave™: In-Memory Database Platforms, Q3 2015

by Noel Yuhanna
August 3, 2015

Why Read This Report

The in-memory database platform represents a new space within the broader data management market. Enterprise architecture (EA) professionals invest in in-memory database platforms to support real-time analytics and extreme transactions in the face of unpredictable mobile, Internet of Things (IoT), and web workloads. Application developers use them to build new applications that deliver performance and responsiveness at the fastest possible speed. Forrester identified the 11 most significant software providers — Aerospike, DataStax, IBM, Kognitio, MemSQL, Microsoft, Oracle, SAP, Starcounter, Teradata, and VoltDB — in the category and researched, analyzed, and scored them against 19 criteria. This report details how well each vendor fulfills Forrester's criteria and where the vendors stand in relation to each other to meet next-generation real-time data requirements.

Key Takeaways

SAP, Oracle, IBM, Microsoft, And Teradata Lead The Pack
Forrester's research uncovered a market in which SAP, Oracle, IBM, Microsoft, and Teradata lead the pack. MemSQL, Kognitio, VoltDB, DataStax, Aerospike, and Starcounter offer competitive options.

The In-Memory Database Platform Market Is New But Growing Rapidly
The in-memory database market is new but growing fast as more enterprise architecture professionals see in-memory as a way to address their top data management challenges, especially to support low-latency access to critical data for transactional or analytical workloads.

Scale, Performance, And Innovation Distinguish The In-Memory Database Leaders
The Leaders we identified offer high-performance, scalable, secure, and flexible in-memory database solutions. The Strong Performers have turned up the heat as high as it will go on the incumbent Leaders, with innovations that many customers find compelling.

The Forrester Wave™: In-Memory Database Platforms, Q3 2015

by Noel Yuhanna
with Leslie Owens, Elizabeth Cullen, and Diane Lynch
August 3, 2015

Table Of Contents

Notes & Resources

Forrester conducted executive strategy briefings in April 2015 and interviewed 11 vendor and user companies: Aerospike, DataStax, IBM, Kognitio, MemSQL, Microsoft, Oracle, SAP, Starcounter, Teradata, and VoltDB.

Related Research Documents

Brief: In-Memory Data Platform Is No Longer An Option — It's A Necessity!

Market Overview: In-Memory Data Platforms

TechRadar™: Big Data, Q3 2014

FORRESTER®

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
+1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com

In-Memory Database Is Gaining Momentum Across All Use Cases

Every second counts when you want to deliver real-time recommendations to customers based on their location, activity, or status. The traditional approach of storing data on disk and later integrating and analyzing it isn't good enough anymore; EA professionals need to deliver the perfect recommendation in seconds, not days or weeks. Storing and processing customer data, events, and clickstreams in memory supports such sub-second, low-latency access. For example, real-time reservations applications for hotels, concerts, sports events, restaurants, and car rentals use in-memory data platforms to deliver enhanced customer experiences that allow users to choose a table, seat, or room of their liking in real time, even when hundreds and thousands of users are doing the same at that moment. Until recently, however, storing and processing larger amounts of data in memory was not an option because it was prohibitively expensive. Today's in-memory database platforms are changing the way we build and deliver systems of engagement and are transforming the practice of analytics, predictive modeling, and business transaction management.

Forrester defines an in-memory database platform as:

A database technology that stores all or partial data in memory on either a single or distributed server to support transactional, operational, and/or analytical workloads.

In-Memory Delivers Extreme Performance And Scalability

Today, in-memory supports many use cases, including analytics, extreme transactions, real-time data access, big data applications, mobile applications, predictive analytics, and operational reporting. The top in-memory workloads commonly seen include:

› **Big data applications that need fast data.** Hadoop primarily focuses on distributed batch processing that often slows down delivery of critical data. Enterprise architects can combine Hadoop with in-memory platforms such as Apache Spark or in-memory databases to process and deliver large amounts of data quickly. For example, an energy company might have a real-time network utilization dashboard across various grids, where each grid's data is aggregated across many data points in a Hadoop cluster before being presented in a dashboard using an in-memory data platform.

› **Mobile apps that need integrated data.** Today, data mobilization enables rich interactions and advanced analytics using devices such as tablets, smartphones, and wearables. Mobile application developers demand data from multiple technology stacks and in real time to deliver a 360-degree view of the customer, product, employee, or business. For example, a mobile application might deliver a dashboard that tracks investments from multiple sources in real time. In-memory offers the ability to deliver integrated data that's critical to support such applications.

› **Real-time apps that are critical for business to succeed.** In-memory data platforms can help deliver real-time apps to support operational applications such as stock trading, fraud detection, counterterrorism, patient health monitoring, machine analysis, or earthquake monitoring. These apps require data 24x7 with low-latency access, and even persisting data causes slowdowns and sometimes cannot be accepted. Although many companies have been using real-time apps for decades, such apps previously required extensive application coding and customization to deliver extreme performance.

› **Customer analytics to deliver improved customer experience.** Major retail stores, midsized retailers, and eCommerce sites have started to exploit in-memory data platforms in a big way. Customer data stored and processed in memory help create an opportunity for businesses to upsell and cross-sell new products to a customer based on his or her likes, dislikes, circle of friends, buying patterns, and past orders. To support granular, personalized customer experiences, in-memory data is critical to deliver faster predictive modeling, enable real-time data access, and process big data quickly.

› **IoT applications that can improve operational efficiency.** Today, most manufacturers deal with highly sophisticated machinery to support their plants, whether building an airplane, car, or tire or bottling wine or soda. When a machine goes down, it can cost a manufacturer millions of dollar every hour, and in some cases, every minute. With IoT sensor, streaming, machine-learning, and in-memory technologies, manufacturers are able to track machines every minute, or even every second, to predict if any machine is likely to fail as well as decide what parts or resources might be needed for repairs if a breakdown does occur.

Strong Growth Is Ahead For The In-Memory Database Market

The in-memory database market is new but growing rapidly as enterprise architects look at supporting new and emerging fast data insights, applications, and processes. We can segment the in-memory database market into two main categories:

› **Traditional database vendors that are extending their platforms to support in-memory.** The in-memory market continues to be dominated by large vendors, such as IBM, Microsoft, Oracle, SAP, and Teradata, that offer broad coverage to support most use cases. These vendors continue to invest heavily in in-memory and market their products aggressively.

› **Specialized, pure-play vendors that offer extreme scale and performance.** The specialized vendors, such as Aerospike, DataStax, Kognitio, MemSQL, Starcounter, and VoltDB, offer more integrated solutions, ease of use, and simplicity to speed up platform and application deployment.

In-Memory Database Evaluation Overview

The in-memory market is extremely competitive because it has become a critical category in data management; pure-play in-memory vendors and traditional database vendors are gunning for a piece of this rapidly emerging market. Customers will benefit as the pace of innovation increases and cost of in-memory further declines to support a petabyte-scaled environment.

The Evaluation Focuses On Scale, Performance, Administration, And Use Cases

After examining past research, user need assessments, and vendor and expert interviews, we developed a comprehensive set of evaluation criteria. We evaluated vendors against 19 criteria, which we grouped into three high-level buckets:

› **Current offering.** To assess the breadth and depth of each vendor's in-memory product set, we evaluated each solution's architectural and operational functionality.

› **Strategy.** We reviewed each vendor's strategy to assess how it plans to evolve its in-memory solution to meet emerging customer demands. We also evaluated each vendor's go-to-market approach, commitment, and direction strategies.

› **Market presence.** To establish each in-memory database product's market presence, we evaluated each solution provider's company financials, adoption, and partnerships.

Evaluated Vendors Meet Functional, Architectural, And Market Presence Criteria

Forrester included 11 vendors in the assessment: Aerospike, DataStax, IBM, Kognitio, MemSQL, Microsoft, Oracle, SAP, Starcounter, Teradata, and VoltDB. Each of these vendors has (see Figure 1):

› **An enterprise-class in-memory offering.** Vendors offer the following core in-memory database functional components, tools, and features: 1) supports core in-memory database features and functionality, including high availability, security, performance, scalability, and management; 2) supports data storage for persistence; 3) offers data integrity and consistency; 4) provides native tools or integrates with third-party vendors to support data loading, unloading, administration, security, integration, data quality, archiving, etc.; 5) supports multiple concurrent queries, transactions, reports, or data access patterns; 6) can be deployed on-premises, via public cloud, or both; and 7) data can be accessed using standard connectivity such as SQL, ODBC/JDBC, XML, or representational state transfer (REST).

› **A standalone in-memory database solution.** We are including only in-memory database products that are not technologically tied to any particular applications (such as enterprise resource planning [ERP] or CRM) or particular business intelligence (BI); business performance solution (BPS); predictive analytics; extract, transfer, and load (ETL); or middleware stack and that do not require embedding in other applications. Products must be supported in a standalone environment.

› **A publicly available, in-memory database release.** Each participating player must actively market an enterprise in-memory database solution that is clearly mentioned on its website marketing literature and brochures. The product version included in the evaluation must have been generally available prior to April 1, 2015.

› **A referenceable install base.** There should be 20 or more enterprise customers using the in-memory database product. Customers must span more than one major geographic region.

› **Interest from Forrester clients.** We are including only in-memory database vendors that have been mentioned by customers in at least five Forrester inquiry calls over the past 12 months.

FIGURE 1 Evaluated Vendors: Product Information And Selection Criteria

Vendor	Product	Version
Aerospike	Aerospike	3.54
DataStax	DataStax Enterprise	4.6
IBM	IBM DB2	10.5
Kognitio	Kognitio Analytical Platform	8.1
Microsoft	SQL Server 2014	
MemSQL	MemSQL	3.2
Oracle	Oracle Database Oracle Database In-Memory Oracle TimesTen In-Memory Database	12.1.0.2
SAP	SAP Hana	SPS09
Starcounter	Starcounter in-memory platform	2
Teradata	Teradata Database	15
VoltDB	VoltDB	v5.0

Inclusion criteria

An enterprise-class in-memory offering. Vendors offer the following core in-memory database functional components, tools, and features: 1) supports core in-memory database features and functionality, including high availability, security, performance, scalability, and management; 2) supports data storage for persistence; 3) offers data integrity and consistency; 4) provides native tools or integrates with third-party vendors to support data loading, unloading, administration, security, integration, data quality, archiving, etc.; 5) supports multiple concurrent queries or transactions or reports or data access patterns; 6) can be deployed on-premises, public cloud, or both; and 7) data can be accessed using standard connectivity such as SQL, ODBC/JDBC, XML, or REST.

A publicly available in-memory database release. Each participating player must actively market an enterprise in-memory database solution that is clearly mentioned on its website marketing literature and brochures. The product version included in the evaluation must have been generally available prior to April 1, 2015.

A referenceable install base. There should be 20 or more enterprise customers using the in-memory database product. Customers must span more than one major geographic region.

A standalone in-memory database solution. We are including only in-memory database products that are not technologically tied to any particular applications (such as ERP or CRM) or particular business intelligence (BI), business performance solution (BPS), predictive analytics, ETL, or middleware stack and that do not require embedding in other applications. Products must be supported in a standalone environment.

Interest from Forrester clients. We will include only in-memory database vendors that have been mentioned by customers in at least five Forrester inquiry calls over the past 12 months.
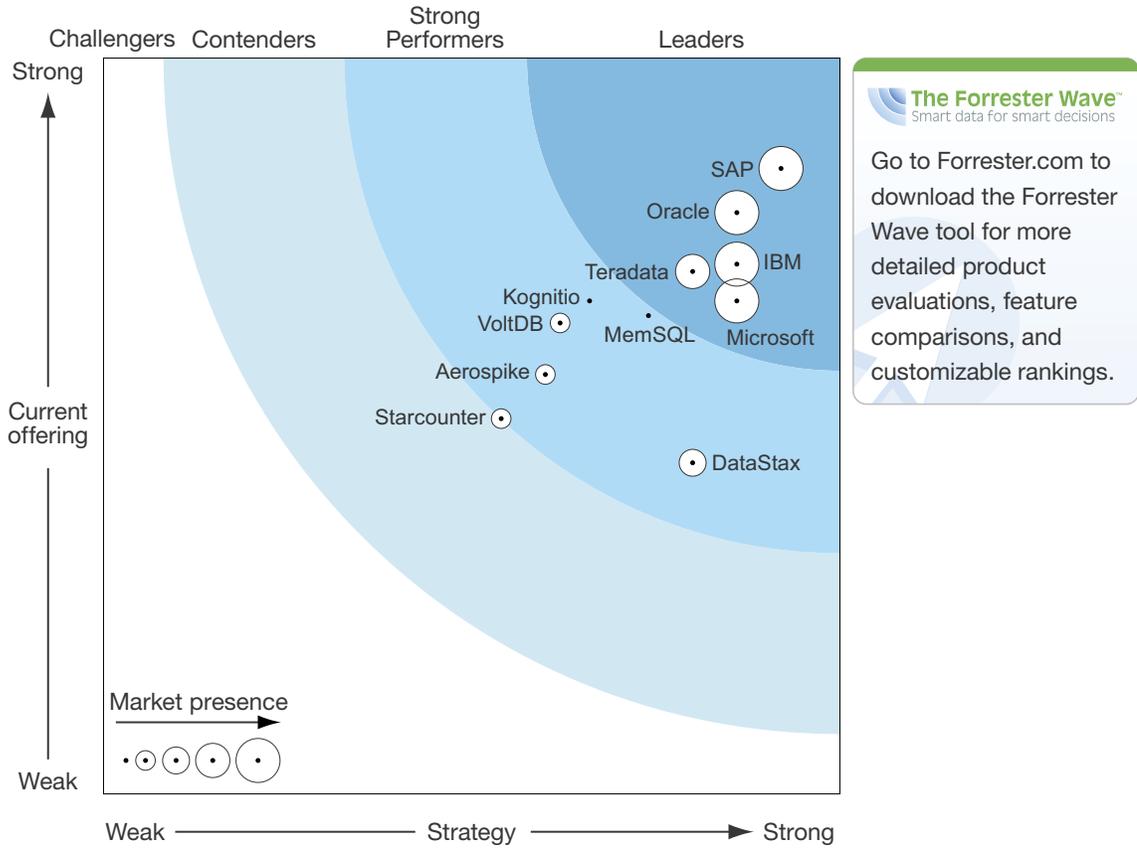
Scale And Broader Use Case Support Distinguish The Leaders

The evaluation uncovered a market in which (see Figure 2):

› **SAP, Oracle, IBM, Microsoft, and Teradata lead with viable enterprisewide solutions.** The in-memory database Leaders — SAP, Oracle, Microsoft, Teradata, and IBM — offer scalable and proven in-memory platforms that support broad workloads and use cases. SAP Hana continues to have the highest adoption, largely because of its distributed and integrated solutions with SAP applications. Oracle offers two in-memory database products — Oracle TimesTen In-Memory Database, a mature, complete in-memory database platform, and Oracle Database In-Memory, which was recently added to extend Oracle Database 12c to enable in-memory support. IBM has ramped up its in-memory offering with its integration of DB2 with IBM DB2 with BLU Acceleration (BLU Acceleration) and IBM PureData System for Analytics to support analytic and operational use cases. Microsoft offers specialized in-memory databases for online analytical processing (OLAP) and online transaction processing (OLTP) workloads, presenting a viable option for Microsoft SQL Server customers. Teradata Intelligent Memory delivers a scalable and integrated in-memory solution that's part of the Teradata Database, used to support real-time analytics, predictive analytics, and extreme operational reporting.

› **MemSQL, Kognitio, VoltDB, DataStax, Aerospike, and Starcounter are competitive.** With innovation that many customers find compelling, this group has turned up the heat as high as it will go on the incumbent Leaders. MemSQL is known for its high-performance transactional capabilities, delivering horizontal scale to support large and complex workloads. Kognitio is known for its row-based, in-memory database platform for analytics, which makes extensive use of a RAM-based system. VoltDB, an open source, in-memory database, continues to have good success offering enterprises a viable alternate to leading database vendor solutions. DataStax uses Apache Cassandra to deliver its in-memory database solution to support operational and analytical use cases. Aerospike continues to be a solid option for enterprises looking for low-latency access to critical data stored in dynamic random-access memory (DRAM), SSD, and Flash. Starcounter offers a viable platform for medium-sized deployments.

This evaluation of the in-memory database market is intended to be a starting point only. We encourage clients to view detailed product evaluations and adapt criteria weightings to fit their individual needs through the Forrester Wave Excel-based vendor comparison tool.

FIGURE 2 The Forrester Wave™: In-Memory Database Platforms, Q3 '15

Go to Forrester.com to download the Forrester Wave tool for more detailed product evaluations, feature comparisons, and customizable rankings.

FIGURE 2 The Forrester Wave™: In-Memory Database Platform, Q3 '15 (Cont.)

	Forrester's weighting	Aerospike	DataStax	IBM	Kognitio	MemSQL	Microsoft	Oracle	SAP	Starcounter	Teradata	VoltDB
CURRENT OFFERING	50%	2.85	2.25	3.60	3.35	3.25	3.35	3.95	4.25	2.55	3.55	3.20
Data management features	10%	2.00	2.00	3.00	3.00	4.00	4.00	4.00	5.00	4.00	3.00	4.00
Scale-out architecture	10%	3.00	2.00	3.00	4.00	4.00	3.00	3.00	3.00	3.00	5.00	3.00
Analytics support	10%	2.00	3.00	4.00	5.00	3.00	3.00	4.00	5.00	2.00	5.00	3.00
Server platform	10%	2.00	2.00	5.00	3.00	2.00	2.00	4.00	4.00	2.00	2.00	2.00
Transaction capabilities	10%	4.00	2.00	4.00	2.00	4.00	3.00	4.00	5.00	3.00	2.00	4.00
Performance and scale	10%	5.00	1.00	3.00	5.00	5.00	4.00	4.00	5.00	5.00	5.00	4.00
High availability and disaster recovery	10%	4.00	4.00	4.00	3.00	3.00	3.00	4.00	4.00	2.00	4.00	4.00
Languages supported	5%	1.00	2.00	4.00	3.00	3.00	4.00	4.00	4.00	2.00	4.00	3.00
Data movement	5%	1.00	3.00	3.00	4.00	2.00	4.00	5.00	3.00	1.00	5.00	2.00
Security	5%	2.00	3.00	3.00	2.00	2.00	3.00	5.00	4.00	1.00	4.00	3.00
Management and monitoring	5%	3.00	1.00	4.00	2.00	2.00	4.00	5.00	4.00	1.00	4.00	2.00
Cloud support	10%	3.00	2.00	3.00	3.00	3.00	4.00	3.00	4.00	2.00	1.00	3.00
STRATEGY	50%	3.00	4.00	4.30	3.30	3.70	4.30	4.30	4.60	2.70	4.00	3.10
Product road map	40%	3.00	4.00	4.00	3.00	4.00	4.00	4.00	4.00	3.00	4.00	4.00
Execution	30%	2.00	5.00	5.00	3.00	3.00	5.00	5.00	5.00	1.00	4.00	1.00
Vision	30%	4.00	3.00	4.00	4.00	4.00	4.00	4.00	5.00	4.00	4.00	4.00
Cost	0%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
MARKET PRESENCE	0%	1.50	2.60	5.00	1.00	1.00	4.50	4.10	5.00	1.50	3.60	1.60
Customers	50%	2.00	3.00	5.00	1.00	1.00	4.00	4.00	5.00	2.00	3.00	2.00
Partners	10%	1.00	3.00	5.00	1.00	1.00	5.00	5.00	5.00	1.00	5.00	2.00
Revenue	40%	1.00	2.00	5.00	1.00	1.00	5.00	4.00	5.00	1.00	4.00	1.00

All scores are based on a scale of 0 (weak) to 5 (strong).

Vendor Profiles

Leaders

› **SAP's razor-sharp focus on in-memory technology is paying off.** SAP Hana is an in-memory, column-oriented, relational database management system (RDBMS) that's designed to support transactions and complex query processing on the same data set. SAP has gained more than 7,000 customers since Hana's launch in 2011.[1] Its go-to-market message focuses on its multinode distributed shared-nothing in-memory platform, real-time analytics, columnar data format, advanced compression, and data services layer. Today, most enterprises use SAP Hana for in-memory data marts, SAP Business Warehouse, SAP S/4 Hana, and SAP Business Suite; additionally, application development with non-SAP applications is starting to ramp up. SAP continues to invest significantly in Hana in engineering and marketing, specifically focusing on dynamic data tiering, big data, granular security, cloud, modeling, and enhanced levels of automation. For customers that want to leverage in-memory on a cloud platform, SAP offers the SAP Hana Cloud Platform, an open platform-as-a-service providing an in-memory database and application services to build, extend, or integrate business applications. Hana is also available on Amazon Web Services (AWS), although most AWS customers are small and medium-sized businesses. SAP Hana might be overkill for customers that are transactional-only or that do not take advantage of the platform's ability to support analytics, calculation, or algorithm processing within its application logic, since SAP Hana is expensive in comparison with other in-memory solutions.

› **Oracle offers multiple in-memory products to support various use cases.** Oracle offers a mature, in-memory database product called Oracle TimesTen In-Memory Database as well as Oracle Database In-Memory, an option that extends Oracle Database 12c to support in-memory. TimesTen runs completely in-memory, requiring no disk access and delivering low-latency access for operational and transactional workloads. Oracle Database In-Memory, on the other hand, requires no change to existing Oracle applications and supports a horizontal scale for OLTP and OLAP applications. With the combination of these two options, customers can practically deploy any type of application — transactional, operational, and analytical. Forrester spoke with several customers that run TimesTen along with Oracle Database In-Memory, supporting a platform for multiple workloads. Oracle's on-premises strategy remains strong; however, the cloud platform (including support for in-memory) is still evolving. Forrester expects it will grow rapidly as Oracle expands its database-as-a-service (DBaaS) offering in the coming years. For customers that need a larger memory footprint, Oracle delivers a scale-out in-memory platform using Oracle Exadata appliance or any Oracle RAC environment. Oracle's road map includes performance enhancements with increased use of vector processing, optimization for scanning and processing data, faster decompression using software-in-silicon, in-memory replication, a high degree of automation, and in-memory extension to Flash. Oracle customers interested in speeding up their transactions or analytical workloads should look at TimesTen and/or consider upgrading to the in-memory option in Oracle Database 12c.

› **IBM delivers a viable in-memory data platform for DB2.** IBM DB2 with BLU Acceleration is an in-memory columnar data platform that speeds up real-time and ad hoc analytical and operational workloads on IBM DB2. IBM's large installed base of database deployments gives it a competitive advantage as it rolls out advanced in-memory capabilities. We see that firms using BLU Acceleration require less data storage because of its cutting-edge data compression technology and can also optimize data access routines without the need of tuning, indexes, and aggregates. BLU Acceleration is tightly integrated with IBM DB2, which means that customers can move to BLU Acceleration from existing DB2 applications without changing any application code. IBM still is ramping up its BLU Acceleration in-the-cloud offering. IBM BLU Acceleration is likely to focus on supporting a broader set of advanced analytical workloads, deeper big data integration, an enhanced cloud solution, greater hardware acceleration, and enhancement in security in the coming years.

› **Microsoft offers specialized in-memory databases for OLTP and OLAP.** Microsoft offers two in-memory database solutions within the SQL Server platform to address OLTP and OLAP workloads. In SQL Server 2014, Microsoft delivered an improved in-memory column-store for data warehousing to support faster business intelligence, analytics, and predictive analytics. In addition, SQL Server 2014 delivers a new in-memory OLTP database platform to support high-performance transactional applications, requiring minimal changes to the application. With both of these solutions, Microsoft SQL Server customers do not have to change any SQL Server applications or queries. Although these solutions are not completely in-memory, customers claim to see significant performance improvements of up to 30 times for OLTP and up to 100 times for OLAP from previous SQL Server releases; the OLAP in-memory offering can transparently leverage disk storage for data exceeding the size of available physical memory. Microsoft's road map will continue to zero in on enhancing the in-memory platform, focusing on broader use cases, including Internet of Things, big data, and real time; handling faster data ingestion; and supporting real-time analytic queries and global operations. Customers running Microsoft SQL Server should look at the SQL Server 2014 in-memory option to improve their transactions and analytics; the upgrade is worth the effort. In addition, Microsoft plans to further enhance its in-memory with improved performance, Microsoft Azure support, automation, and elastic scale capabilities in the upcoming SQL Server 2016 release.

› **Teradata Intelligent Memory delivers a credible power pack for analytics.** Teradata Intelligent Memory is part of the Teradata Database, which keeps active data in memory for faster analytics, insights, and reporting. It uses multi-temperature technology that automatically stores the hottest data in DRAM, warm data in SSD/Flash, and cold data on economical disk drives. Although Teradata does not currently offer a complete in-memory database or support for Intelligent Memory in the cloud, it might in the future, given its mature scale-out data management platform. Teradata's strength lies in delivering high performance and scale for analytical workloads, appliance integration, and strong high availability and disaster recovery, as well as broad integration with tools and programming languages. Teradata's road map includes enhancements to in-memory processing and optimizations to Intelligent Memory. Teradata will continue to apply advanced

engineering techniques to support larger and more complex analytical and operational workloads. Unlike other larger database management system (DBMS) vendors, Teradata does not support comprehensive transactional capabilities today; therefore, enterprises must move data using ETL and replication tools or custom scripts into the Teradata platform.

Strong Performers

› **MemSQL delivers a real-time, focused database for transactions and analytics.** MemSQL is a distributed in-memory database that delivers full atomicity, consistency, isolation, and durability (ACID) compliance and extreme performance to support transactional, operational, and analytical workloads. As a purpose-built database for low-latency access to real-time and historical data, MemSQL uses a familiar SQL interface and a horizontally scalable, distributed architecture that runs on commodity hardware or in the cloud. Although Forrester estimates that MemSQL has only about 50 customers, it continues to execute well on its vision and is expanding its market with some large customer deployments, including Comcast, Shutterstock, and Zynga. We find that several companies use MemSQL to better predict and react to opportunities by extracting value in their data to drive new revenue. MemSQL's road map includes improved development with Apache Spark and geospatial, increasing seamless cloud deployment, and adding management features and analysis functions. Customers that are building new transactional and analytical applications that need extreme performance and low-latency access and want a single database platform should look at MemSQL.

› **Kognitio's massively parallel in-memory analytical platform is powerful.** Kognitio is one of the leading vendors that deliver scalable big data analytics for data science and business intelligence. Its massively parallel in-memory data platform combines SQL and NoSQL and reads big data directly from Hadoop, data warehouses, and cloud while integrating with business applications via industry-standard APIs. Kognitio offers a flexible architecture that can be deployed on-premises, in the cloud, as appliance, or as a service and has deployments across most vertical industries, including retail, media, financial services, and telecommunication. Although Kognitio does not deliver transaction-processing-centric solutions, its core strength lies in supporting complex analytics in real time and near-real time. Kognitio's road map includes in-memory mirroring for resilience and direct JSON loading and processing as well as improvements in scale and performance. Enterprises looking at high-performance, low-latency big data analytical workloads will find Kognitio worth consideration.

› **VoltDB offers a viable alternative for streaming analytics and transactions.** VoltDB is an in-memory database platform that combines streaming analytics with transactions in a single integrated platform. The product is an ACID-compliant and distributed shared-nothing in-memory database developed by several database researchers, including Michael Stonebraker.[2] VoltDB relies on horizontal partitioning of data to scale out on commodity hardware and public clouds. It also supports synchronous replication within the database cluster to support high availability.

Unlike other vendors, VoltDB is available as open source software under the Affero General Public License (AGPL) as well as under a commercial license. It supports multistatement transactions via pre-compiled Java stored procedures. Customers use VoltDB to support real-time analytics and low-latency transactional and operational applications across the financial services, energy, advertising, security, and gaming industries. VoltDB's road map includes enhancing rapid data ingestion, real-time analytics processing, active/active replication across data centers, enhancing export capabilities with support for multiple export targets, geospatial functions that are accessible with SQL, and MPP import connectors. Enterprises should shortlist VoltDB when building new low-latency applications where every second counts.

› **DataStax delivers a viable option, combining performance and availability.** DataStax is a database platform based on Apache Cassandra, built for performance and availability for Internet of Things (IoT), web, mobile, and extreme-scale applications. With more than 400 customers, DataStax continues to execute well to support global applications that demand low-latency access to critical information. Although DataStax cannot run complete in-memory, its masterless, shared-nothing, peer-to-peer architecture scales across data centers to deliver a strong, distributed global data platform for applications that demand 24x7 availability and scale. DataStax's road map includes increasing the addressable in-memory option and automated tiered storage to support DRAM, SSD, and disk automatically. We expect DataStax to support a complete in-memory data platform in the coming years, leveraging its core distributed data platform. For customers that need a combination of extreme performance and availability, DataStax is one vendor to shortlist.

› **Aerospike's NoSQL in-memory solution delivers extreme performance and scale.** Aerospike is an in-memory database that supports real-time, context-driven applications. Developers use Aerospike, an open source, Flash-optimized, in-memory NoSQL key-value store, for caching, as a context store, and to simplify scaling with smaller clusters. Aerospike is written in C and operates in three layers: a Flash-optimized data layer, a self-managed distribution layer, and a cluster-aware client layer. The data layer is optimized to store data in solid-state drives, RAM, or traditional rotational media. Top use cases for Aerospike include fraud detection, real-time bidding, personalized web portals, eCommerce search, and context-driven applications. One customer runs more than 4.5 million transactions per second with more than 100 terabytes of data replicating across eight clusters in four data centers. However, customers are limited by the fact it runs only on Linux; does not support ODBC, JDBC, or simple object access protocol (SOAP); and has limited third-party vendor tooling. Aerospike's road map includes focusing on cloud, security, API integration, scale, and broader use cases. Overall, Aerospike is a great solution for customers that need speed and want to integrate the solution using software development kits (SDKs) with Java, .NET, and other programming frameworks.

› **Starcounter delivers a viable in-memory NoSQL database.** Starcounter, founded in 2006, is a fully ACID-compliant, in-memory NoSQL database that is often used to support high-performance transactional and real-time applications. Customers use Starcounter to build web applications, mobile services, and real-time applications in categories such as gaming, advertising, and telecom.

It's optimized for throughput, concurrency, and low-latency access. It also features tight integration with web standards such as HTTP, REST, and JSON. Forrester estimates that there are fewer than 150 customers using Starcounter; however, some are large corporations around the globe. Starcounter's road map includes extending the platform to support IoT, ease of development, improved performance, and a better user interface. Although Starcounter doesn't have all the bells and whistles when compared with larger in-memory vendors, such as IBM, SAP, or Oracle, it offers a strong, economical platform, especially when supporting a NoSQL data store.

Engage With An Analyst

Gain greater confidence in your decisions by working with Forrester thought leaders to apply our research to your specific business and technology initiatives.

Analyst Inquiry

Ask a question related to our research; a Forrester analyst will help you put it into practice and take the next step. Schedule a 30-minute phone session with the analyst or opt for a response via email.

Learn more about inquiry, including tips for getting the most out of your discussion.

Analyst Advisory

Put research into practice with in-depth analysis of your specific business and technology challenges. Engagements include custom advisory calls, strategy days, workshops, speeches, and webinars.

Learn about interactive advisory sessions and how we can support your initiatives.

Supplemental Material

Online Resource

The online version of Figure 2 is an Excel-based vendor comparison tool that provides detailed product evaluations and customizable rankings.

Data Sources Used In This Forrester Wave

Forrester used a combination of three data sources to assess the strengths and weaknesses of each solution:

› **Vendor surveys.** Forrester surveyed vendors on their capabilities as they relate to the evaluation criteria.

› **Executive Strategy Briefing.** Forrester asked participants to describe the company's background, positioning, value proposition, customer base, and strategic vision. We used findings from these calls to validate details of each vendor's strategy and vision.

› **Customer reference calls.** To validate product and vendor qualifications, Forrester also fielded an online survey to six of each vendor's current customers.

The Forrester Wave Methodology

We conduct primary research to develop a list of vendors that meet our criteria to be evaluated in this market. From that initial pool of vendors, we then narrow our final list. We choose these vendors based on: 1) product fit; 2) customer success; and 3) Forrester client demand. We eliminate vendors that have limited customer references and products that don't fit the scope of our evaluation.

After examining past research, user need assessments, and vendor and expert interviews, we develop the initial evaluation criteria. To evaluate the vendors and their products against our set of criteria, we gather details of product qualifications through a combination of lab evaluations, questionnaires, demos, and/or discussions with client references. We send evaluations to the vendors for their review, and we adjust the evaluations to provide the most accurate view of vendor offerings and strategies.

We set default weightings to reflect our analysis of the needs of large user companies — and/or other scenarios as outlined in the Forrester Wave document — and then score the vendors based on a clearly defined scale. These default weightings are intended only as a starting point, and we encourage readers to adapt the weightings to fit their individual needs through the Excel-based tool. The final scores generate the graphical depiction of the market based on current offering, strategy, and market presence. Forrester intends to update vendor evaluations regularly as product capabilities and vendor strategies evolve. For more information on the methodology that every Forrester Wave follows, go to http://www.forrester.com/marketing/policies/forrester-wave-methodology.html.

Integrity Policy

All of Forrester's research, including Forrester Wave evaluations, is conducted according to our Integrity Policy. For more information, go to http://www.forrester.com/marketing/policies/integrity-policy.html.

Endnotes

[1] Source: "SAP Announces Second Quarter and First Half 2015 Results," SAP press release, July 21, 2015 (http://global.sap.com/corporate-en/investors/newsandreports/news.epx?articleID=24735&category=45).

[2] Michael Stonebraker is an adjunct professor at the Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (MIT CSAIL), where he is also cofounder and codirector of the Intel Science and Technology Center for Big Data. Dr. Stonebraker was a professor of computer science at the University of California at Berkeley for 29 years and has won numerous industry and academic awards. He has been a database pioneer since the early 1970s, when he developed relational DBs Ingres. He also architected Illustra, Vertica, object-relational Postgres, in-memory DB VoltDB, and Mariposa, a massively distributed federated database system commercialized through Cohera Corp. Most recently, Stonebraker's work has focused on specialized database management systems, such as SciDB for scientific data. In total, he has founded or cofounded nine startups, the most recent being Tamr, where he is chief technology officer (CTO). Tamr's products take on the problem of curation of big data at enterprise scale. See the "An Interview With The 2014 Turing Award Winner: Dr. Michael Stonebraker" Forrester report.

We work with business and technology leaders to develop customer-obsessed strategies that drive growth.

PRODUCTS AND SERVICES

› Core research and tools
› Data and analytics
› Peer collaboration
› Analyst engagement
› Consulting
› Events

Forrester's research and insights are tailored to your role and critical business initiatives.

ROLES WE SERVE

Marketing & Strategy Professionals
CMO
B2B Marketing
B2C Marketing
Customer Experience
Customer Insights
eBusiness & Channel Strategy

Technology Management Professionals
CIO
Application Development & Delivery
› Enterprise Architecture
Infrastructure & Operations
Security & Risk
Sourcing & Vendor Management

Technology Industry Professionals
Analyst Relations

CLIENT SUPPORT

For information on hard-copy or electronic reprints, please contact Client Support at +1 866-367-7378, +1 617-613-5730, or clientsupport@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.