



Transform Your Applications For Real Time

FEATURING RESEARCH FROM FORRESTER

The Forrester Wave™: Translytical Data Platforms, Q4 2017

Transform Your Applications For Real Time

TRANSLYTICAL DATA PLATFORMS BECOME CRITICAL FOR THE INSTANT EXPERIENCE

Applications today must respond instantly and intelligently for your business to stay competitive. This requires a different type of data platform – one that delivers instant access to business data and supports multiple workloads and use cases. Implementing a data platform that meets these requirements comes with several associated costs and risks. The right platform not only needs to meet performance requirements, but it also needs to incorporate the most up-to-date methods of solving data-related problems. Innovation is enhanced when the platform is cost-efficient in its delivery, requires the fewest resources to maintain, and is versatile enough to be used across a range of modern scenarios.

“Translytical data platforms can support many types of use cases, including real-time insights, machine learning, streaming analytics, and extreme transactional processing.”, says this Forrester Wave, “The sweet spot is the ability to perform all of these workloads within a single database..” This is really the crux that makes Translytical Data Platforms the right choice for modern application architectures.

CHOOSING AMONG THE RIGHT TRANSLYTICAL DATA PLATFORMS

Among the criteria used to select between translytical platforms, each vendor evaluated offers the following core translytical functional components, tools, and features:

- 1) A core data platform capability, including full support for high availability, security, performance, scalability, and management
- 2) Data logging for persistence
- 3) Data integrity and consistency to deliver trusted data
- 4) Native tools or integration with third-party vendors to support loading, unloading, administration, security, integration, data quality, archiving and other data management functions
- 5) Support for concurrent queries, transactions, reports, and data access patterns
- 6) On-premises or public cloud deployments, or both
- 7) Access to data using standard connectivity such as SQL, ODBC/JDBC, XML, or REST.

IN THIS DOCUMENT

1 Transform Your Applications For Real Time

5 The Forrester Wave™: Translytical Data Platforms, Q4 2017

22 About Redis Enterprise

The Forrester Wave™: Translytical Data Platforms, Q4 2017

The 12 Providers That Matter Most And How They Stack Up

by Noel Yuhanna and Mike Gualtieri

November 28, 2017

Why Read This Report

Analytics at the speed of transactions has become an important agenda item for organizations. Translytical data platforms, an emerging technology, deliver faster access to business data to support various workloads and use cases. EA pros can use them to drive new business initiatives. Forrester identified the 12 most significant translytical vendors — Aerospike, DataStax, GigaSpaces, IBM, MemSQL, Microsoft, NuoDB, Oracle, Redis, SAP, Splice Machine, and VoltDB — and researched, analyzed, and scored them against 25 criteria.

Key Takeaways

Twelve Translytical Vendors Compete In This Growing Market

Among the translytical vendors Forrester evaluated, we found five Leaders, four Strong Performers, and three Contenders.

Enterprise Architects Are Looking To Support Larger Translytical Deployments

The translytical data platform market is growing because more enterprise architecture (EA) pros see translytical as critical for their enterprise data strategy.

Performance, Scale, Security, And Use-Case Support Are Key Differentiators

The Leaders we identified support a broader set of use cases and more integrated translytical capabilities. The Strong Performers have turned up the heat on the incumbent Leaders to offer more data management features, and Contenders are ramping up their offerings to support more use cases and higher performance.

The Forrester Wave™: Translytical Data Platforms, Q4 2017

The 12 Providers That Matter Most And How They Stack Up

by [Noel Yuhanna](#) and [Mike Gualtieri](#)

with [Gene Leganza](#) and Jun Lee

November 28, 2017

Table Of Contents

The Translytical Data Platform Creates New Business Opportunities

Translytical Platforms Support Demanding Use Cases In A Streamlined Architecture

Translytical Data Platform Evaluation Overview

Evaluated Vendors And Inclusion Criteria

Vendor Profiles

Leaders

Strong Performers

Contenders

Supplemental Material

Related Research Documents

[Emerging Technology: Translytical Databases Deliver Analytics At The Speed Of Transactions](#)

[The Forrester Wave™: Big Data Fabric, Q4 2016](#)

[The Forrester Wave™: Big Data Warehouse, Q2 2017](#)



SHARE REPORTS WITH COLLEAGUES. ENHANCE YOUR MEMBERSHIP WITH RESEARCH SHARE.

FORRESTER

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

FORRESTER

© 2017 Forrester Research, Inc. Opinions reflect judgment at the time and are subject to change. Forrester®, Technographics®, Forrester Wave, TechRadar, and Total Economic Impact are trademarks of Forrester Research, Inc. All other trademarks are the property of their respective owners. Forrester Research, Inc. is not responsible for any content or copyright infringement on this page. For more information, contact [forrester.com](#) or +1 866-367-7378

The Translytical Data Platform Creates New Business Opportunities

We find that enterprises usually have separate databases, data warehouses, and/or data lakes to support varying workloads, such as transactional, operational, and analytical. These data silos and technology stacks make it difficult for EA pros to support real-time insights and analytics without incessantly moving data around. Data must move from transactional systems to operational systems and then to analytical systems, which slows down processing, integration, and timely insights. Stale reports, missing data, perishable insights, and the lack of advanced real-time analytics create an unbearable state for any enterprise that needs fresh insights to remain competitive.¹ Disparate technology stacks compromise timely, integrated data delivery to business users, customers, and partners.

TRANSLYTICAL PLATFORMS SUPPORT DEMANDING USE CASES IN A STREAMLINED ARCHITECTURE

Translytical data platforms can support many types of use cases, including real-time insights, machine learning, streaming analytics, and extreme transactional processing. The sweet spot is the ability to perform all of these workloads within a single database. For example, businesses can store and process customer data in a single integrated translytical platform, enabling them to upsell and cross-sell new products based on customer likes, dislikes, buying patterns, friend circles, and past orders. The top translytical database workloads Forrester has seen are:

- › **Real-time apps — where every second counts.** Translytical can help deliver real-time apps to support operational applications, such as stock trading, fraud detection, counterterrorism, patient health monitoring, machine analysis, and earthquake monitoring. These apps require data 24x7 with low-latency access, and even persisting data can cause unacceptable slowdowns. While many companies have been using real-time apps for decades, long before translytical databases, such apps required extensive application design, coding, and customization.
- › **Internet-of-things (IoT) analytics on operational data.** Today, most manufacturers rely on highly sophisticated machinery to support their plants. When a machine goes down, it can cost a manufacturer millions of dollars every hour — and in some cases, every minute. With IoT sensors, streaming, machine learning, and in-memory technologies, manufacturers can track machines every second to predict likely failures as well as to decide what parts or resources they might need for repairs if a breakdown does occur.
- › **Connected data apps — where integrated business data is critical.** Traditional extract, transform, and load (ETL) processes fail to deliver real-time changes. Translytical platforms overcome this challenge by delivering a real-time, trusted view of critical business data, ensuring that the source of information is accurate to guarantee consistency across the organization. For example, a customer's address might be stored on five or more different databases, and a change by one application might not be visible to other app users right away. In this case, storing all customer-critical data in-memory in a translytical database allows all business applications to use it, delivering consistency and integrity.

- › **Continuous learning.** As enterprises increasingly rely on machine learning models to make predictions about customers, business processes, and operations, they need to keep these models fresh with new data. Translytical databases support ongoing training, retraining, and monitoring of machine learning models without the need to move data to external machine learning platforms, which is costly and time-consuming.

Translytical Data Platform Evaluation Overview

To assess the state of the translytical data platform market and see how the vendors stack up against each other, Forrester evaluated the strengths and weaknesses of top translytical data platform vendors. After examining past research and Forrester Wave evaluations, user need assessments, and vendor and expert interviews, we developed a comprehensive set of 25 evaluation criteria, which we grouped into three high-level buckets:

- › **Current offering.** To assess the breadth and depth of each vendor's translytical data platform offering, we evaluated each solution's architectural and operational features and functionality.
- › **Strategy.** We reviewed each vendor's strategy to assess how it plans to evolve its translytical data platform services to meet emerging customer demands. We also evaluated each vendor's go-to-market approach, commitment, and direction strategies.
- › **Market presence.** To establish each translytical data platform offering's market presence, we evaluated each provider's company financials, adoption, and partnerships.

EVALUATED VENDORS AND INCLUSION CRITERIA

Forrester included 12 vendors in the assessment: Aerospike, DataStax, GigaSpaces, IBM, MemSQL, Microsoft, NuoDB, Oracle, Redis, SAP, Splice Machine, and VoltDB. Each of these vendors has (see Figure 1):

- › **An enterprise-class translytical offering.** Vendors offer the following core translytical functional components, tools, and features: 1) a core data platform capability, including full support for high availability, security, performance, scalability, and management; 2) data logging for persistence; 3) data integrity and consistency to deliver trusted data; 4) native tools or integration with third-party vendors to support loading, unloading, administration, security, integration, data quality, archiving, and other data management functions; 5) support for concurrent queries, transactions, reports, and data access patterns; 6) on-premises or public cloud deployments, or both; and 7) access to data using standard connectivity such as SQL, ODBC/JDBC, XML, or REST.
- › **A standalone translytical solution.** We included only products that are not technologically tied to any particular applications (such as packaged applications or software-as-a-service offerings) or particular business intelligence (BI), business performance solution, predictive analytics, ETL, or middleware stack and that do not require embedding in other applications. Products must be supported in a standalone environment to support various workloads.

- › **A publicly available release.** Each participating vendor must actively market an enterprise translytical data platform or similar solution. The product version included in the evaluation must have been generally available as either software or a cloud service as of July 1, 2017.
- › **A referenceable install base.** Evaluated vendors must have at least five enterprise customers using their translytical product or service. In addition, customers must span more than one major geographic region, including the US.
- › **Customer interest.** Forrester included only vendors that customers mentioned several times during inquiry calls in the past 12 months related to translytical and/or related topics.
- › **Client inquiries and/or technologies that put the vendor on Forrester's radar.** Forrester clients often discuss the vendors and products through inquiries and interviews; alternatively, the vendor may, in Forrester's judgment, warrant inclusion or exclusion in this evaluation because of technology trends and market presence.

FIGURE 1 Evaluated Vendors: Product Information And Selection Criteria

Company	Product name	Version name
Aerospike	Aerospike Database	3.14
DataStax	DataStax Enterprise	DataStax Enterprise 5.1 (includes DSE Studio 2.0 and DSE OpsCenter 6.0)
GigaSpaces	InsightEdge Platform; Platform Components: 1) XAP, 2) InsightEdge Analytics	12.2
IBM	IBM Db2 IBM Db2 for z/OS with IBM Db2 Analytics Accelerator	11.1.2.2 IBM Db2 for z/OS = 12, IBM Db2 Analytics Accelerator = 5.1
MemSQL	MemSQL	5.8.4
Microsoft	SQL Server 2016	
NuoDB	NuoDB	NuoDB 2.6
Oracle	Oracle Database	12c Release 2
Redis	Redis Enterprise	4.5
SAP	SAP HANA, SAP Data Hub, SAP Agile Data Preparation	SAP HANA 2.0 SPS01, SAP Data Hub 1.0, SAP Agile Data Preparation 1.0 SP16
Splice Machine	Splice Machine	2.6
VoltDB	VoltDB Enterprise Edition	7.5.1

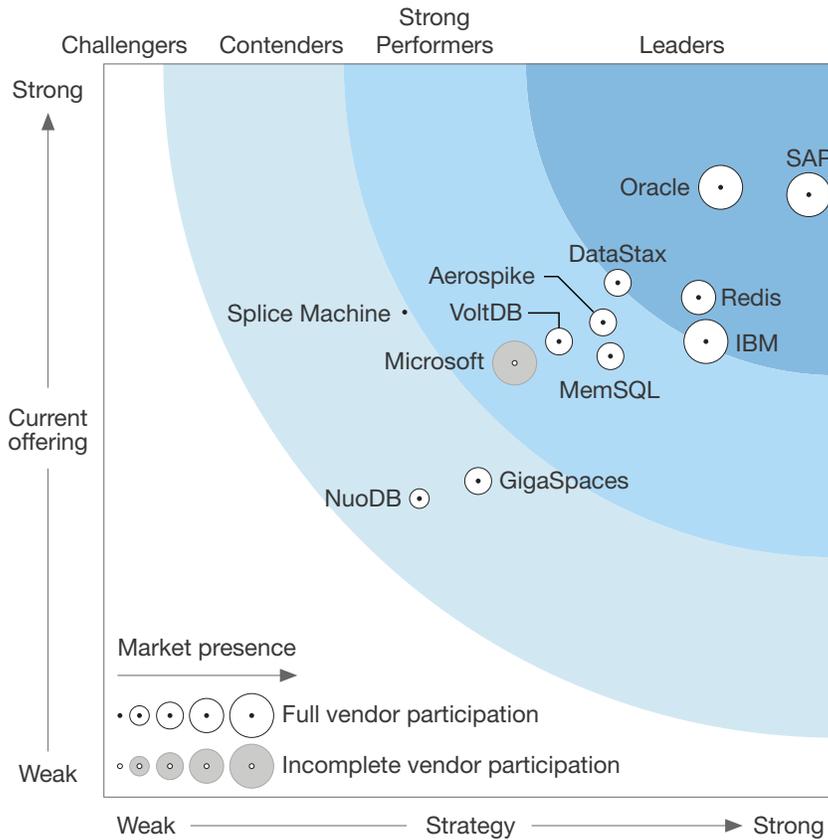
FIGURE 1 Evaluated Vendors: Product Information And Selection Criteria (Cont.)**Vendor inclusion criteria**

- **An enterprise-class translytical offering.** Vendors offer the following core translytical functional components, tools, and features: 1) a core data platform capability, including full support for high availability, security, performance, scalability, and management; 2) data logging for persistence; 3) data integrity and consistency to deliver trusted data; 4) native tools or integration with third-party vendors to support loading, unloading, and other data management functions; 5) support for concurrent queries, transactions, reports, and data access patterns; 6) on-premises or public cloud deployment, or both; and 7) access to data using standard connectivity such as SQL, ODBC/JDBC, XML, or REST.
- **A standalone translytical solution.** We included only products not technologically tied to any particular applications (such as packaged applications or software-as-a-service offerings) or business intelligence; business performance solution; predictive analytics; extract, transform, and load (ETL); or middleware stack. Products must not require embedding in other application; they must be supported in a standalone environment to support various workloads.
- **A publicly available release.** Vendors must actively market an enterprise translytical data platform or similar solution. The product version included in Forrester's evaluation must have been generally available as either software or a cloud service as of July 1, 2017.
- **A referenceable install base.** Evaluated vendors must have at least five enterprise customers using the translytical product or service. In addition, customers must span more than one major geographic region, including the US.
- **Customer interest.** Forrester included only vendors that customers mentioned several times during Forrester inquiry calls in the past 12 months related to translytical and/or related topics.
- **Client inquiries and/or technologies that put the vendor on Forrester's radar.** Forrester clients often discuss the vendors and products through inquiries and interviews; alternatively, the vendor may, in Forrester's judgment, warrant inclusion or exclusion in this evaluation because of technology trends and market presence.

Vendor Profiles

This evaluation of the translytical data platform 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 (see Figure 2).

FIGURE 2 Forrester Wave™: Translytical Data Platforms, Q4 '17



FORRESTER RESEARCH
The Forrester Wave™

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

FIGURE 2 Forrester Wave™: Translytical Data Platforms, Q4 '17 (Cont.)

	Forrester's weighting	Aerospike	DataStax	GigaSpaces	IBM	MemSQL	Microsoft	NuoDB	Oracle	Redis	SAP	Splice Machine	VoltDB
Current Offering	50%	3.23	3.50	2.15	3.10	3.00	2.95	2.03	4.15	3.40	4.10	3.30	3.10
Development	50%	2.55	2.50	2.40	3.70	3.00	3.90	1.65	4.50	2.90	5.00	2.70	2.70
Deployment	50%	3.90	4.50	1.90	2.50	3.00	2.00	2.40	3.80	3.90	3.20	3.90	3.50
Strategy	50%	3.40	3.50	2.55	4.10	3.45	2.80	2.15	4.20	4.05	4.80	2.05	3.10
Ability to execute	35%	3.00	3.00	3.00	5.00	4.00	3.00	2.00	5.00	4.00	5.00	2.00	3.00
Road map	40%	4.00	3.00	2.00	4.00	3.00	3.00	3.00	4.00	4.00	5.00	2.00	3.00
Community	10%	3.00	5.00	1.00	3.00	1.00	1.00	1.00	1.00	3.00	3.00	1.00	1.00
Implementation support	15%	3.00	5.00	4.00	3.00	5.00	3.00	1.00	5.00	5.00	5.00	3.00	5.00
Market Presence	0%	2.25	3.00	2.60	4.50	2.80	4.50	1.25	4.75	3.30	5.00	0.80	2.95
Product revenue	30%	2.00	3.00	2.00	5.00	2.00	5.00	1.00	5.00	2.00	5.00	0.00	2.00
Install base	30%	2.00	3.00	4.00	5.00	3.00	5.00	1.00	5.00	5.00	5.00	0.00	3.00
Market awareness	25%	3.00	3.00	2.00	3.00	4.00	3.00	2.00	4.00	3.00	5.00	2.00	4.00
Partnerships	15%	2.00	3.00	2.00	5.00	2.00	5.00	1.00	5.00	3.00	5.00	2.00	3.00

LEADERS

- › **SAP crushes translytical workloads.** SAP HANA is a shared-nothing, in-memory data platform, the core of SAP's translytical platform, which supports many use cases, including real-time applications, analytics, translytical apps, systems of insight, and advanced analytics. Enterprises use the platform for in-memory data marts, SAP Business Warehouse, and applications like SAP S/4HANA, and SAP Business Suite. SAP Vora extends the translytical platform by supporting the Apache Spark execution framework to deliver enriched interactive analytics on Hadoop. SAP HANA is also available on many public clouds, such as Amazon Web Services (AWS) and SAP Cloud Platform; although most AWS customers are small and medium-sized businesses, that's starting to change. SAP's translytical platform might be overkill for organizations that don't have large translytical requirements because some enterprise buyers perceive it to be more expensive than less capable solutions.

- › **Oracle leaves no transaction or analytics stone unturned.** Oracle offers Oracle Database In-Memory, an option that extends Oracle Database 12c to support both transactions and analytics in the same database. It requires no change to existing Oracle applications and supports horizontal scale for online transaction processing (OLTP) and online analytical processing (OLAP) applications. For larger-memory footprints, the vendor offers its Oracle Exadata appliance. Its road map includes performance enhancements with faster decompression using software-in-silicon, in-memory replication, a high degree of automation, in-memory flash extension, and NVRAM support. Customers like its ease of use, performance, and general technology and security capabilities.
- › **Redis stretches open source to support translytical workloads.** Redis offers an enterprise in-memory database (based on open source) that supports transactional, analytical, and operational workloads. It supports both relaxed and strong consistency depending on the application requirement, with a flexible and schemaless model that's built on a shared-nothing symmetric architecture. Today, Redis Enterprise runs both on-premises and on AWS, Google Cloud, IBM SoftLayer, and Microsoft Azure, and it is integrated into several platform-as-a-service environments. Customers use Redis Enterprise to support real-time analytics, translytical, data ingestion, social media, customer intelligence, operational reporting, and 360-degree-view-of-the-customer initiatives. Customers like its business value, support for broad use cases, performance, and customer support.
- › **IBM has multiple translytical database offerings.** IBM's Db2 with BLU Acceleration and Informix with Informix Warehouse Accelerator deliver transactional and analytical workloads on a variety of distributed platforms, leveraging an in-memory, columnar approach that requires no application changes or tuning. In addition, IBM Db2 for z/OS integrates with the IBM Db2 Analytics Accelerator (IDAA) to enable real-time analytics on transactional data. Enterprises using BLU Acceleration often require less data storage because of its data compression technology and the ability to optimize data access patterns without the need of tuning, indexes, and aggregates. Customers like its ease of use, business capabilities, security, and overall total cost of ownership.
- › **DataStax Enterprise is a data platform based on Apache Cassandra.** DataStax Enterprise (DSE) supports various workloads, including IoT, mobile, transactional, operational, and extreme-scale applications. DSE extends and embeds Apache Spark to natively support operational analytics within the platform. With more than 500 customers, DataStax continues to execute well to support global applications that demand low-latency access to critical data. DSE has a masterless, shared-nothing architecture with in-memory capabilities and built-in analytics/search that can run on-premises and in the cloud to support various workloads. DataStax's road map includes increasing the addressable in-memory option and automated tiered storage to support DRAM, solid-state drives, and disk automatically. Customers like its scalability, lower cost, customer support, and distributed architecture.

STRONG PERFORMERS

- › **Aerospike is built for hybrid-memory architecture.** Aerospike Database supports various applications, including real-time systems of engagement, operational intelligence, and advanced analytics. It is a flash-optimized, in-memory, NoSQL, key-value store database that simplifies the development and deployment of applications needing low-latency access. Top use cases include fraud detection, real-time bidding, personalized web portals, eCommerce search, and context-driven applications. One large customer runs more than 4.5 million transactions per second, with over 100 terabytes of data replicating across eight clusters in four data centers for a mission-critical, customer-facing application. Customers like its ability to drive business innovation, its ease of use, and its leveraging of Flash storage.
- › **MemSQL delivers a viable translytical platform.** MemSQL Enterprise is a distributed in-memory database that delivers full ACID compliance and performance to support transactional, operational, and analytical workloads. As a purpose-built database for low-latency data access, MemSQL uses SQL and a horizontally scalable, distributed architecture that runs on commodity hardware or the public cloud. MemSQL continues to execute well on its vision and is expanding its market share with some large deployments, including Akamai Technologies, Dell EMC, Kellogg, Pinterest, and Samsung. MemSQL's road map includes improvements with Apache Spark; enabling further integration with cloud platforms; and improving performance, scale, and administration for larger in-memory deployments.
- › **VoltDB leverages its shared-nothing, in-memory architecture for translytics.** VoltDB's in-memory database combines streaming analytics with transactions in a single integrated platform. The product is an ACID-compliant, immediately consistent, distributed, shared-nothing in-memory database. VoltDB relies on horizontal partitioning of data to scale out on commodity hardware, on-premises, or in the public cloud. Unlike other vendors, VoltDB is available as open source software under the Affero General Public License (AGPL) as well as under a commercial license. Customers use VoltDB to support real-time analytics and low-latency transactional and operational applications across the telecom, financial services, energy, and advertising industries.
- › **Microsoft ramps a unified solution for OLTP and OLAP workloads.** Microsoft offers two in-memory database solutions within the SQL Server platform to address translytical workloads. In SQL Server 2016, Microsoft supports an in-memory column store for data warehousing to support faster BI, analytics, and predictive analytics, as well as an in-memory OLTP feature to support high-performance transactional applications. Although these solutions are not completely in-memory, customers claim to see significant performance improvements. Microsoft's road map focuses on improving its platform, broader use cases (including the IoT, big data, and real time), faster data ingestion, and supporting global applications. Microsoft did not fully participate in this evaluation.

CONTENDERS

- › **Splice Machine offers a scalable translytical platform.** Splice Machine is an OLAP platform that seamlessly integrates a scale-out SQL relational database management system (RDBMS) with machine learning, streaming, and notebooks to power predictive applications. Splice Machine meets the unique demands of predictive applications through OLAP and machine learning to learn from the past as well as streaming and traditional ACID-compliant OLTP capabilities to plan and act in real time. Examples include predicting and planning around shortages in supply chains, avoiding code-blue events in hospitals, and predicting and planning maintenance for IoT devices. Splice Machine is open source and available on-premises or as a fully managed cloud service (i.e., DBaaS).²
- › **GigaSpaces emphasizes transactions.** If you can make it on Wall Street, you can make it anywhere. GigaSpaces was the first vendor to create an in-memory data grid for Java that is used by investment banks and other industries. With GigaSpaces InsightEdge Platform and its core engine, XAP, enterprises can deliver real-time analytics and insights for big data scenarios. Today, many organizations in financial services, retail, transportation, telco, and healthcare leverage GigaSpaces to support instant insights, extreme transactions, low-latency operational intelligence, and advanced analytics workloads.
- › **NuoDB delivers clustered price/performance but is new to analytics.** NuoDB combines scale-out capabilities, elasticity, and continuous availability with transactional consistency and durability to support hybrid cloud applications. NuoDB's distributed database appears as a single logical SQL database to the application, offering independent scale related to transaction processing and storage management. NuoDB's peer-to-peer architecture enables near real-time analytics on high-throughput transaction processing requirements. It supports active-active clustering, memory architecture, API-driven administration, and on-demand scale-out and scale-in capabilities. Customers like its simplicity model, automation, commitment to improving scale and performance with every release, and the business value it creates.

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

To help you put research into practice, connect with an analyst to discuss your questions in a 30-minute phone session — or opt for a response via email.

[Learn more.](#)

Analyst Advisory

Translate research into action by working with an analyst on a specific engagement in the form of custom strategy sessions, workshops, or speeches.

[Learn more.](#)

Webinar

Join our online sessions on the latest research affecting your business. Each call includes analyst Q&A and slides and is available on-demand.

[Learn more.](#)



Forrester's research apps for iOS and Android.

Stay ahead of your competition no matter where you are.

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. Click the link at the beginning of this report on Forrester.com to download the tool.

DATA SOURCES USED IN THIS FORRESTER WAVE

Forrester used a combination of three data sources to assess the strengths and weaknesses of each solution. We evaluated the vendors participating in this Forrester Wave, in part, using materials that they provided to us by August 8, 2017.

- › **Vendor surveys.** Forrester surveyed vendors on their capabilities as they relate to the evaluation criteria. Once we analyzed the completed vendor surveys, we conducted vendor calls where necessary to gather details of vendor qualifications.

- › **Customer reference online survey.** To validate product and vendor qualifications, Forrester also fielded an online survey with at least two of each vendor's current customers.

THE FORRESTER WAVE METHODOLOGY

We conduct primary research to develop a list of vendors that meet our criteria for evaluation in this market. From that initial pool of vendors, we 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 evaluation — and then score the vendors based on a clearly defined scale. We intend these default weightings to serve only as a starting point and 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

We conduct all our research, including Forrester Wave evaluations, in accordance with our Integrity Policy. For more information, go to <http://www.forrester.com/marketing/policies/integrity-policy.html>.

Endnotes

¹ See the Forrester report “[Perishable Insights — Stop Wasting Money On Unactionable Analytics.](#)”

² DBaaS: database-as-a-service.

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.



ABOUT REDIS ENTERPRISE

What makes Redis Enterprise a unique data platform? Redis Enterprise nails the basics of a translytical data platform – the highest performance with the least resources, a range of use-cases all delivered with uncompromising performance and modularly extensible to every data processing scenario in a modular fashion. But what truly differentiates is:

- **Active-active** distributed datasets with CRDTs (Conflict free replicated datatypes) make it ideal for a microservices or geographically distributed architectures. With the loosely coupled architectures prevalent in microservices-based environments, with instances that are spun up and torn down independently, distributed shared datasets are the norm. The simplest, most effective way to gain consistent views across shared datasets without increasing latency or sacrificing throughput are CRDTs. Mathematically proven to converge to the same correct state, [CRDT-based active-active](#) databases allow for elastic, flexible loose-coupling demanded by microservices or by geographically distributed applications.
- **Built-in high performance search** and secondary indexing. Don't confuse this with search functionality in other data platforms. [Benchmarked](#) to be 5x faster than market leading search engines, the scale and performance of simultaneous indexing and search accomplished within Redis Enterprise is peerless.
- **Intelligent tiered access to DRAM and SSDs.** Redis Enterprise can store and process data across a combination of memory and flash storage, optimizing for the lowest latencies (<1ms) and highest throughput, through smart distribution of the hottest data to memory and cooler data to slower but cost effective, Flash-based SSDs. The resulting infrastructure can be over [80% more cost effective](#) than other options for similar data throughput and latencies.