

OFFERING OVERVIEW

There Is a HeatWave for Databases: Oracle MySQL Database Service Merges OLTP and OLAP

Oracle MySQL Database Service With HeatWave Offers the Best of Both Worlds in a Breathtaking, Elegant Architecture



Holger Mueller
Vice President and Principal Analyst

TABLE OF CONTENTS

Executive Summary3

About Oracle MySQL Database Service With HeatWave4

Functional Capabilities.....9

Analysis and Observations 19

Recommendations 21

Related Research 23

Endnotes 26

Analyst Bio..... 27

About Constellation Research 28



EXECUTIVE SUMMARY

This Offering Overview examines Oracle MySQL Database Service with HeatWave. The report describes the underlying market trends, introduces the vendor, and presents key differentiators for Oracle's offering. It continues with an analysis of strengths and weaknesses of the vendor and concludes with a set of tangible and actionable recommendations for CxOs.

Business Themes



New C-Suite



Data to Decisions



Technology Optimization

ABOUT ORACLE MYSQL DATABASE SERVICE WITH HEATWAVE

Overview

MySQL was created in 1995 by David Axmark and Michael Widenius, named after Widenius' oldest daughter My, and released in 1996 (see Figure 1). MySQL quickly created a large market share inside of enterprises as a solid and stable SQL transactional database. Widenius sold MySQL AB to Sun Microsystems in 2008, which itself was acquired by Oracle in 2009.¹

Oracle has been a good steward for MySQL, supporting and enhancing the MySQL platform since the acquisition and adding new features and capabilities, such as a high-availability (HA) version. The latest innovation has been the addition of HeatWave in December 2020, a cloud-native in-memory query accelerator designed to accelerate OLAP and other complex queries within Oracle MySQL Database Service.

Figure 1. MySQL Founders David Axmark and Michael Widenius (Left) and Geir Høydalsvik, Oracle's Senior Software Development Director for MySQL (Right)



Source: Wikipedia

Market Definition

Since the 1950s, databases have remained not only central to computing but also a foundational layer for enterprise software. The persistent nature of enterprise software requires that the information captured must be available after the user's session ends. Databases ensure and deliver this persistence.

Moreover, databases must allow software and users to access and process information; how they accomplish this is a key area of differentiation between database providers. At their core, databases organize a collection of data objects, including schemas, tables, queries, reports, and views.

With the advent of the cloud, deployment options now reflect the shift of IT loads from on-premises to the cloud. As a result, the availability of next-gen databases in the public cloud will become more attractive as enterprises shift to next-generation applications.

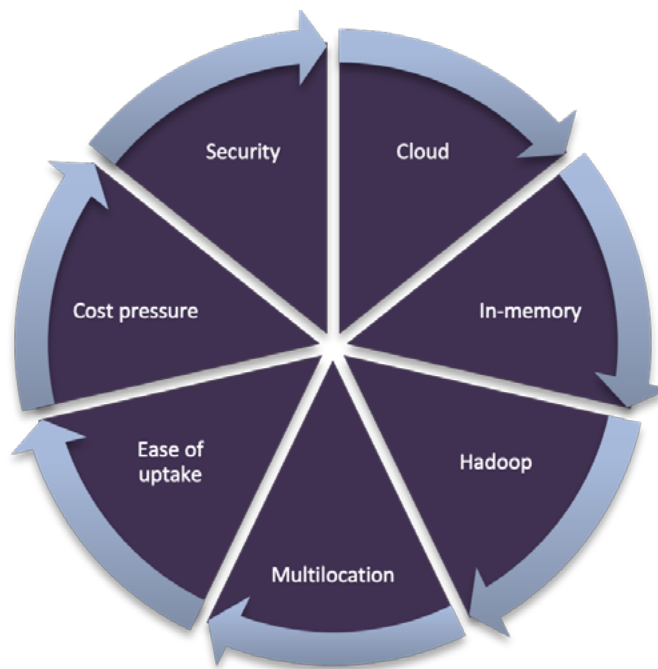
Constellation's conversations with clients show us that database choices matter. Why? Switching databases is hard and expensive, and doing so often poses a risk to business continuity. Leaders must make wise decisions about databases that provide a long-term return on investment, reduce overall operating costs, and deliver on enterprise agility.

Market Trends

Constellation has identified seven key market trends for next-generation databases (see Figure 2):

- 1. Enterprises expect cloud deployment options.** Customers expect their vendors to have a cloud deployment option. How this is delivered does not matter. For example, offering cloud-native options, using third-party infrastructure as a service (IaaS), delivering on vendor-run cloud infrastructure, or even relying on the infrastructure of a partner are all valid options. Remarkably, all vendors—except for those offering the two cloud-native options—make their databases available as on-premises products. Notably, the Hadoop-centric database vendors have been the slowest to offer public cloud deployments, relative to the start date of their first commercial offerings. CxOs expect elasticity of cloud resources, consumption-based pricing, and flexibility of licensing models.

Figure 2. Seven Market Trends Defining Next-Gen Databases



Source: Constellation Research

- 2. New designs support in-memory computing.** Moore's Law has not only made hard disk drives (HDDs) cheaper but also lowered the cost of random-access memory (RAM). The result: Enterprises can run large parts or, in some cases, the complete transactional enterprise resource planning (ERP) databases of large multinational enterprises in memory. Performance benefits include the ability to process information in real time. Traditional vendors keep a subset of data in (expensive) memory. Consequently, how memory is managed, accessed, and consumed is where the vendors differ. Hadoop-based vendors use Apache Spark and Hive. Traditional RDBMS vendors use memory sparingly and in an organic fashion, moving data to an in-memory database when beneficial for system performance or placed in memory by system administrators. In-memory vendors take a more radical approach and place the entire system in memory at all times. Meanwhile, cloud-based vendors, given the novelty of their offering, have not shared their in-memory philosophy; however, given the use cases, Constellation expects the uptake of in-memory options to be similar to that of traditional RDBMS vendors.
- 3. Options must coexist and integrate with Hadoop.** Hadoop has profoundly changed the data storage and retrieval market while massively transforming enterprise best practices for analytics. For the first time, enterprises can store all of their electronic information in one place, without

knowing what they want to analyze and remaining commercially viable. In addition, a data storage and retrieval architecture can bridge data silos.

4. **Globalization and regulation increase requirements for multiple locations.** Enterprises are more global than in the past and must act globally more than ever. With limitations in performance and statutory requirements regarding data residency, database vendors can no longer pass the responsibility to the customer to solve this challenge. Customers expect solution offerings to support multiple locations where their databases can be operated, either through their own infrastructure or a partner's.
5. **Implementation plans assume ease of deployment.** Enterprises need to move faster and cannot afford to be slowed down. The speed, ease of deployment, and time to go live are key differentiators for database vendors. Critical success factors include helping customers to use a new product effectively, partnering closely with customers during first implementations, and making their product easier to deploy. The application of automation based on machine learning (ML) and artificial intelligence (AI) will improve implementation success.
6. **Open source options increase cost pressures on everyone else.** With the rise of mostly open source, cloud-based databases, the cost pressure on the other vendors has increased. Vendors see themselves forced to justify the license costs of their traditional databases when the base version of an open source-based database is free. And the revenue potential and wallet share for the add-ons that can be licensed in the open source market are not the same as for the traditionally expensive database licenses. All traditional database vendors can already see the pressure on their license prices. Constellation expects this trend to continue.
7. **Security remains paramount to clients.** Enterprise-grade use requires a set of security capabilities to prevent data breaches, cyberattacks, and ransomware to avoid potential liability challenges for an enterprise. But enterprises are also ready to experiment and evaluate when only limited security concepts are available; Hadoop databases are an example. In the end, the value of insight gleaned from data often trumps security. However, large-scale and production deliveries require an acceptable set of security capabilities. To an increasing degree, enterprises expect security capabilities to be automatically deployed. Finding the right balance between insight and security is key for CxOs.

The OLTP vs. OLAP Chasm

Since the release of the very first databases, there has been a continuous diverging demand between online transactional processing (OLTP) and online analytical processing (OLAP), which has led to numerous specialized offerings in each of these areas. It has been widely accepted that the diverging needs of transactions and analytics demand different databases, with an extract, transform, and load (ETL) layer acting as the connector, mostly from OLTP to OLAP.

Figure 3 compares the different characteristics of OLTP and OLAP.

Figure 3. OLTP vs. OLAP

| FEATURE | ONLINE TRANSACTIONAL PROCESSING (OLTP) | ONLINE ANALYTICAL PROCESSING (OLAP) |
|------------------|---|--|
| Purpose | Control and run essential business operations in real time, and reports in batch mode | Plan, solve problems, support decisions, discover hidden insights |
| Query nature | Simple SQL queries | (Typically) complex queries |
| Basic operations | INSERT, UPDATE, DELETE | SELECT |
| Response time | Milliseconds | As long as it takes for value of analytics |
| Source | Transactions | Aggregated data from transactions for analytical purposes |
| Data updates | Short, fast updates initiated by users, programs, and interfaces | Periodical refresh via ETL |
| User roles | Customer-facing personnel, clerks, online shoppers | Knowledge workers such as data analysts, business analysts, and executives |
| Database design | Normalized databases for efficiency | Denormalized databases for analysis speed |
| Examples | MySQL, IBM Db2, Oracle Database, SAP ASE, AWS RDS | IBM Cognos, MicroStrategy, Microsoft Power BI, Oracle Business Intelligence Enterprise Edition, Tableau, Google BigQuery, AWS Redshift |

Source: Constellation Research

It has been widely accepted that OLTP databases are bad at OLAP and OLAP databases are bad at OLTP. The reality has been the coexistence of different databases, often offered by different vendors and typically connected by ETL tools from another set of vendors.

The interesting fact is that both database types typically are accessed via one common programming and query language—SQL. So, it has not been programmatic requirements that have led to the two database categories, but the very nature of their operational requirements.

The consequence for enterprises has been higher costs to operate multiple databases and manage the connection across them. But beyond the operational cost, there has been an information and strategic cost, with the analysis always being out of date. The chase for the fabled “insight to action” was largely hamstrung by the OLTP/OLAP status quo.

So, what if OLTP and OLAP could be bridged for MySQL applications? This is what a team at Oracle Labs set out to investigate.

FUNCTIONAL CAPABILITIES

Best of Both OLTP and OLAP in MySQL Database Service With HeatWave

The Oracle Labs team that set out to bridge the OLTP versus OLAP chasm had to work within the following constraints:

- **Preserve Standard SQL as access language.** Standard (unaltered) SQL had to remain the common programming language for both transactional and analytical workloads.
- **Keep it 100% backward-compatible.** No code changes required for running existing MySQL workloads with the new architecture. This would ensure that the vibrant third-party ecosystem of ISVs could operate their offerings seamlessly, with no code changes or testing.

- **Achieve extreme performance.** Performance of analytic workloads was critical while not compromising transactional performance.
- **Optimize for a commodity cloud.** The system has to be designed for the cloud—to achieve cloud scalability, use commodity cloud services (interconnect, storage, compute) to obtain best performance at the lowest cost.
- **Machine learning-based automation.** Workload-driven machine learning intelligently automates operations such as provisioning and other services.

The solution was to extend the code of the MySQL optimizer. For all queries that come to MySQL, the MySQL optimizer makes a cost-based decision on whether it would be faster to execute the query with the InnoDB engine or with HeatWave, the new in-memory query accelerator. For most of the analytic and complex transactional queries, HeatWave is the faster choice.

Oracle was able to create this very elegant solution because it owns and delivers the code of the MySQL optimizer. With these code changes, the team achieved the seamless operational requirement that allows existing MySQL applications to run on the new MySQL offering as well as the seamless operation of ISV applications.

Equally important was Oracle's ability to make the HeatWave engine scale really well. Using a massively parallel architecture, HeatWave achieved the design goal of 400x faster performance for 400GB of data than standard MySQL (see Figure 4) that could be running on-premises or any cloud. As the data size increases, the performance difference increases beyond 400x.

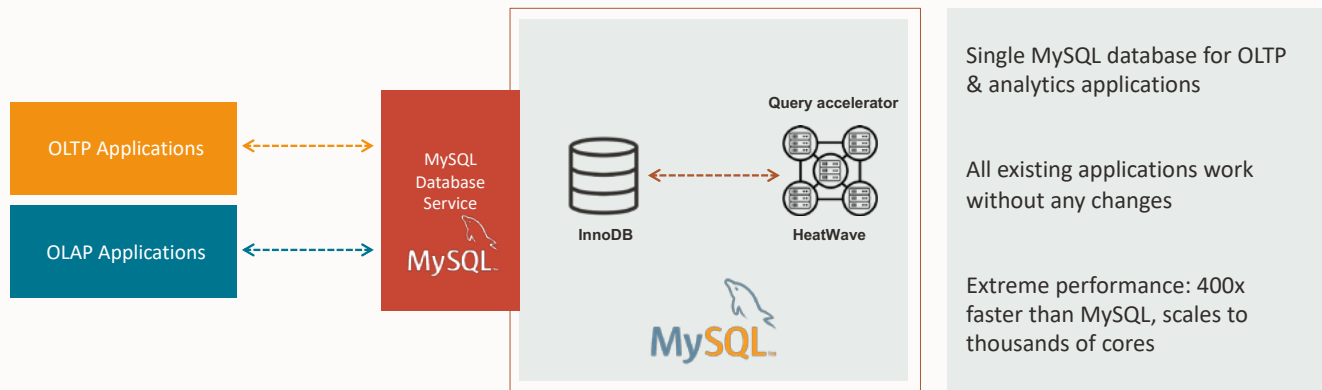
Backward Compatibility Designed From the Beginning

For a new database to succeed, it is critical to attract customers and users. A brand-new database with a new programming language and porting costs for existing workloads—whether those workloads come from enterprises or ISVs—will see only a slow rate of growth. It takes time to rewrite code, whether you are an enterprise or an ISV.

Figure 4. Oracle MySQL HeatWave Functionality

MySQL HeatWave

The *only* MySQL service with a massively-scalable native analytics engine



Source: Oracle


This is why it was critical that HeatWave not require any syntax/code changes. As a result of the architecture chosen, all existing workloads continue to run with Oracle MySQL Database Service with HeatWave (see Figure 5):

- **OLTP applications.** Regardless of the nature of the OLTP application—from software as a service (SaaS) to fintech workloads, social, or e-commerce applications—they all continue to run with zero code changes and therefore zero need for testing or redeployment.
- **OLAP application.** Workloads from OLAP applications such as Oracle Analytics Cloud and Tableau run seamlessly, with no code changes.

HeatWave's Superelegant Architecture Combines OLTP and OLAP

Oracle chose a superelegant architectural approach to make sure the analytics capability addition of HeatWave keeps Oracle MySQL fully backward-compatible with existing workloads. Oracle was able to achieve this because it owns and maintains the underlying source code of MySQL. By

Figure 5. How Oracle MySQL Database Service With HeatWave Enables Existing OLTP and OLAP Workloads—With No Code Changes Required



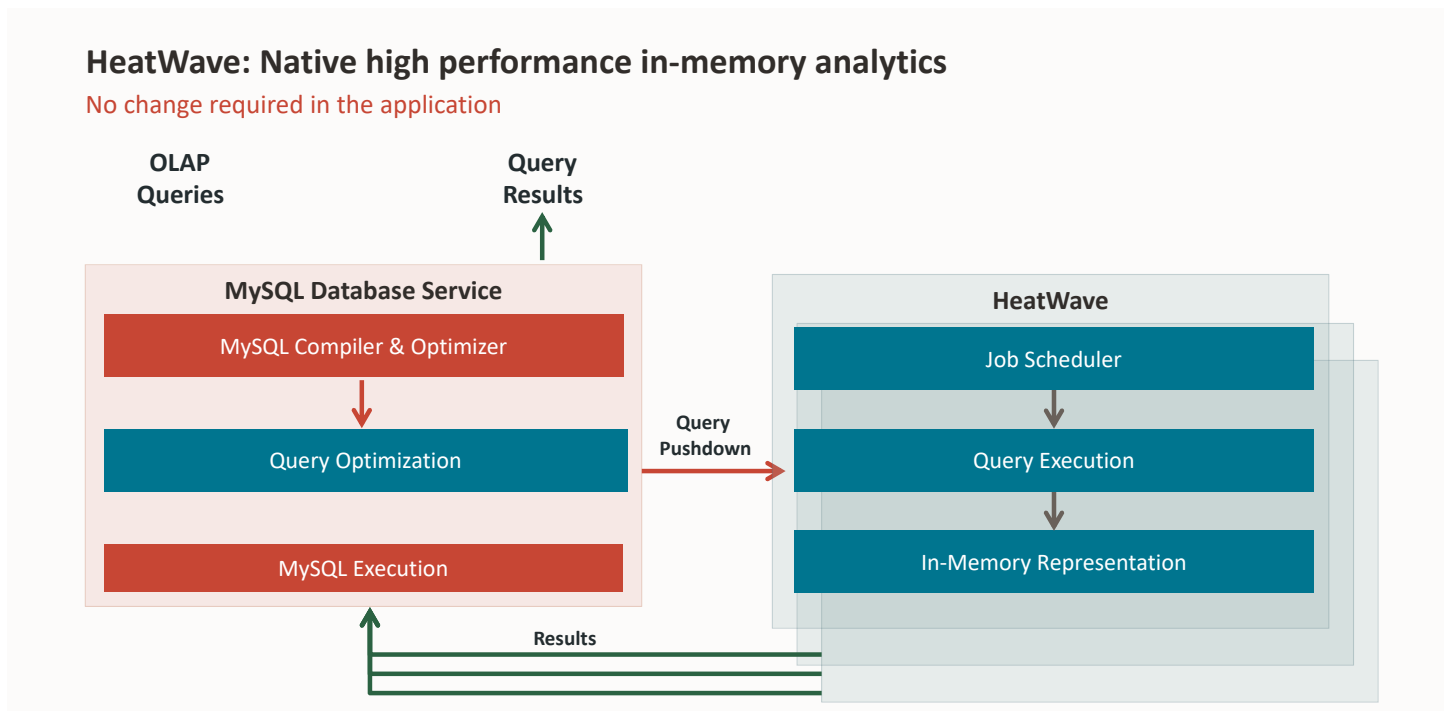
Source: Oracle

changing the code in the Oracle MySQL optimizer, Oracle was able to redirect queries that would be faster with the new engine to HeatWave and the rest continue to be executed with the traditional MySQL engine.

Specifically, it works as follows (see Figure 6):

1. A query arrives to the MySQL database service.
2. The MySQL optimizer evaluates the cost of executing this query with the traditional MySQL engine versus the HeatWave engine.
3. If the cost is cheaper with HeatWave, the MySQL optimizer pushes the query to the HeatWave engine.

Figure 6. Query Pushdown for High-Speed In-Memory Processing



Source: Oracle

4. Inside the HeatWave engine, a new HeatWave optimizer further optimizes the query, comes up with the physical execution plan, and passes it on to the HeatWave execution engine, which uses an in-memory representation of the data. The HeatWave engine is massively scalable and can run in parallel up to 64 servers.
5. The HeatWave engine returns the results to the MySQL execution engine, which presents the results to the application as if they were achieved by the MySQL Database Service.

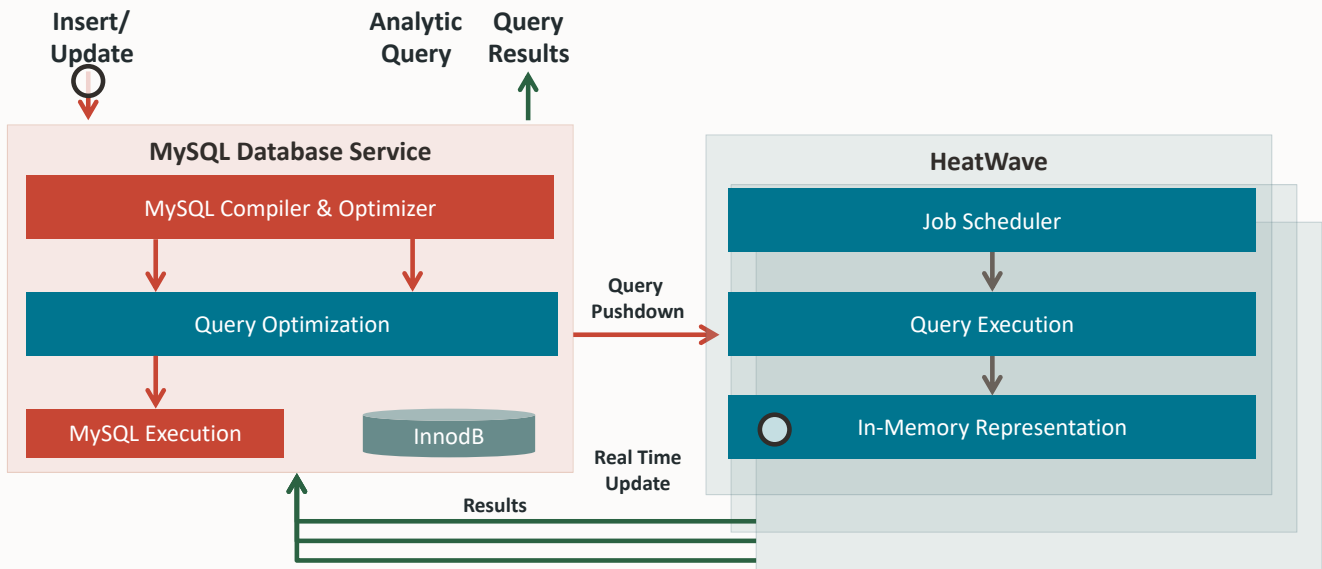
Practically, Oracle has “embedded” a highly scalable and performant query accelerator into its MySQL Database Service with the HeatWave engine.

One of the key challenges in existing solutions where MySQL customers run analytics with a different database is that the data in the analytic system is out of sync with the data in the MySQL database. The HeatWave team came up with an architecture that can accommodate transactional changes in an efficient and time-critical way as well (see Figure 7).

Figure 7. Insight to Action Becomes Feasible With Real-Time Analytics

HeatWave: Analytics in real time

Data inserted into MySQL available for analytic queries in real time



Source: Oracle

When an insert/update command hits the Oracle MySQL Database Service, those changes are persisted in InnoDB, which is the storage engine for the MySQL database. At the same time, these changes are propagated in real time to the in-memory representation of the HeatWave engine. As a result, any query that is issued to the MySQL database following the insert/update statement always sees the latest changes regardless of whether the query is executed by the MySQL engine or the HeatWave engine.

The resulting architecture is elegant, allowing the prioritization of queries to enable the Holy Grail of insight to action for enterprise decision-making. When enterprises can decide in subsecond real time, they have the opportunity to make better decisions—and when the results of these decisions trigger more transactions, then enterprises can improve decision-making again, based on the new transactional reality. This opens up a whole new set of enterprise automation capabilities and processes that enable what matters most: Enterprise Acceleration.²

Seamlessly Offload Analytics Workloads to the Cloud

A long-standing challenge for IT has been the sizing of compute infrastructure to run workloads optimally. Traditionally, with a focus on transactional workloads on-premises, the analytic workloads got short shrift. Often enterprises would not have the capacity to appropriately handle the growing computing needs of analytical workloads.

The challenge of sizing analytical workloads is due to their large storage needs and massive but infrequent compute needs. In the ideal world, analytical workloads would run in the cloud, which is characterized by its architectural and commercial elasticity.

Unfortunately, few analytical workloads have been able to move to the cloud—the main reasons being the lack of on-premises-to-public-cloud ETL tools, network costs, and the need to build a custom solution.

Oracle MySQL has had a long-term option to use built-in replication to move transactions across instances, whether those instances were deployed on-premises or in the public cloud. With the availability of MySQL Database Service with HeatWave in Oracle Cloud Infrastructure (OCI), all main obstacles to moving analytical workloads to the public cloud have been addressed (see Figure 8).

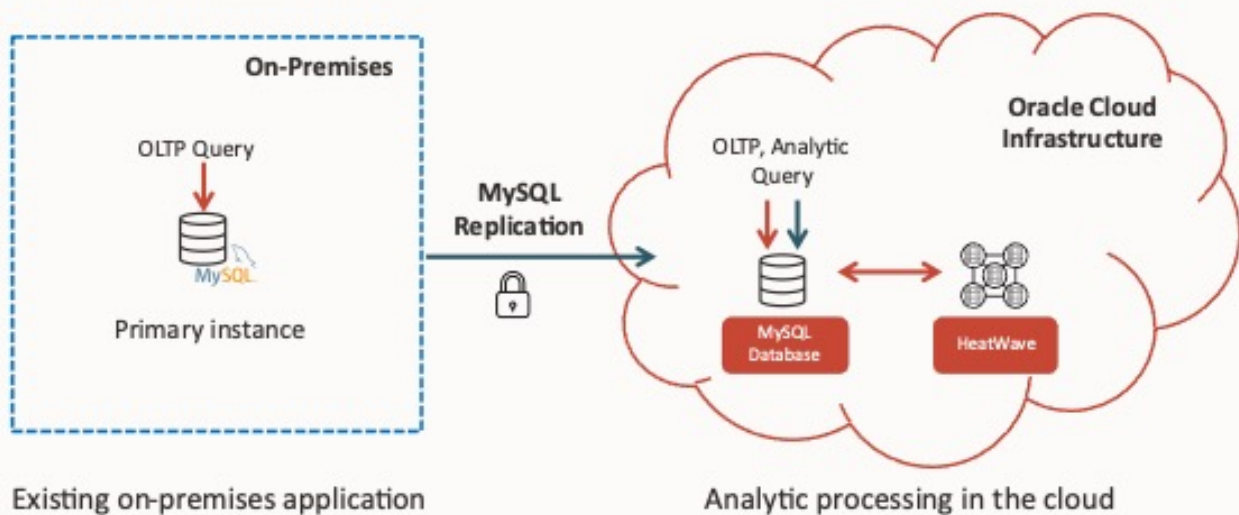
Oracle MySQL provides a proven, standardized, and product-supported replication mechanism for moving transactional data to the cloud, addressing both the ETL and custom setup concerns. The network latency remains, of course, but the benefits of having a cloud-scale platform for analytical workloads should outweigh that latency consideration with its insight and decision velocity/quality benefits.

A Massive Scale and Performance Architecture Powers HeatWave

Analytical workloads are some of the most demanding workloads that enterprises run. They require specialized architectures that need to be high-performing while at the same time as cost-efficient and effective as possible to keep the cost per insight at a minimum.

Figure 8. Analytics for Both Cloud and On-Premises

Easily run analytics against your on-premises MySQL databases



Source: Oracle

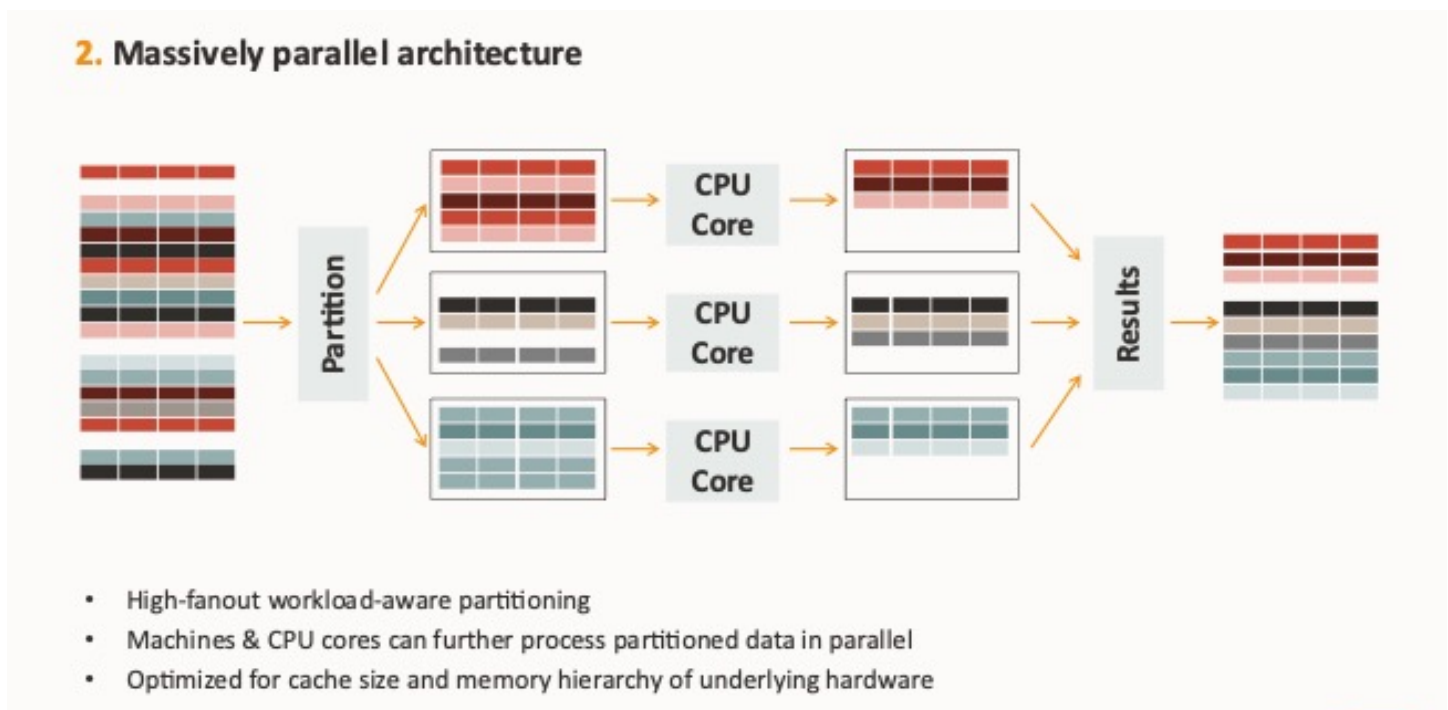
HeatWave achieves both performance and cost goals with the following capabilities:

- 1. Cost efficiencies by leveraging OCI.** HeatWave uses the most cost-effective commodity services from OCI including object storage, commodity interconnect, and compute, and benefits from OCI's high availability, elasticity, and security as well as from the economies of scale. It's priced less than OCI compute instances, so talk about value for your money—HeatWave is off the charts.
- 2. Use of cost-efficient computing elements.** HeatWave uses OCI shapes that provide the lowest cost per TB of memory and have optimized the software stack to achieve a great balance of compute, memory, and network processing.
- 3. In-memory hybrid columnar processing for maximum performance.** HeatWave uses columnar processing to provide the performance analytical workloads demand. It achieves the performance via vectorized execution of the relevant columns, which can be broken into different vectors and multiple chunks as needed. HeatWave manages the vectors in chunks, which benefit from multicore scalability inside of a single partition. Beyond that, HeatWave scales horizontally across n partitions.

4. **A massively parallelized architecture guarantees performance.** HeatWave achieves massive parallelism via massive partitioning, assigning workloads to as many CPU cores as needed for the analytical workload, and then combining back the results. The partitioning architecture is designed for high-fanout workloads and can be done at near memory bandwidth; the machines and CPUs used can further partition the data in parallel (see #3 above), and with Oracle designing the entire architecture, HeatWave can optimize partition sizes based on the cache size of the hardware (see Figure 9).

5. **The “chip-to-click” benefits at work—algorithms distributed for OCI.** Oracle has achieved tremendous performance and TCO gains from a vertically integrated stack—from the silicon all the way through to the user’s click (we refer to this as the Oracle chip-to-click stack)—and it delivers in spades for HeatWave as well. Specifically, HeatWave partitions the data to fit into the cache size, processes the partitions as fast as possible—for example, by using hand-tuned primitives such as Advanced Vector Extensions 2 (AVX2) for OCI CPUs—and optimizes the network for OCI interconnects by infusing intelligence into the scheduling of execute and transfer workloads.

Figure 9. HeatWave’s Massively Parallel Architecture



Source: Oracle

This helps HeatWave achieve excellent scalability across a large number of servers while using commodity hardware and services.

Oracle MySQL Database Service as a First-Class Citizen With Other Oracle Cloud Services

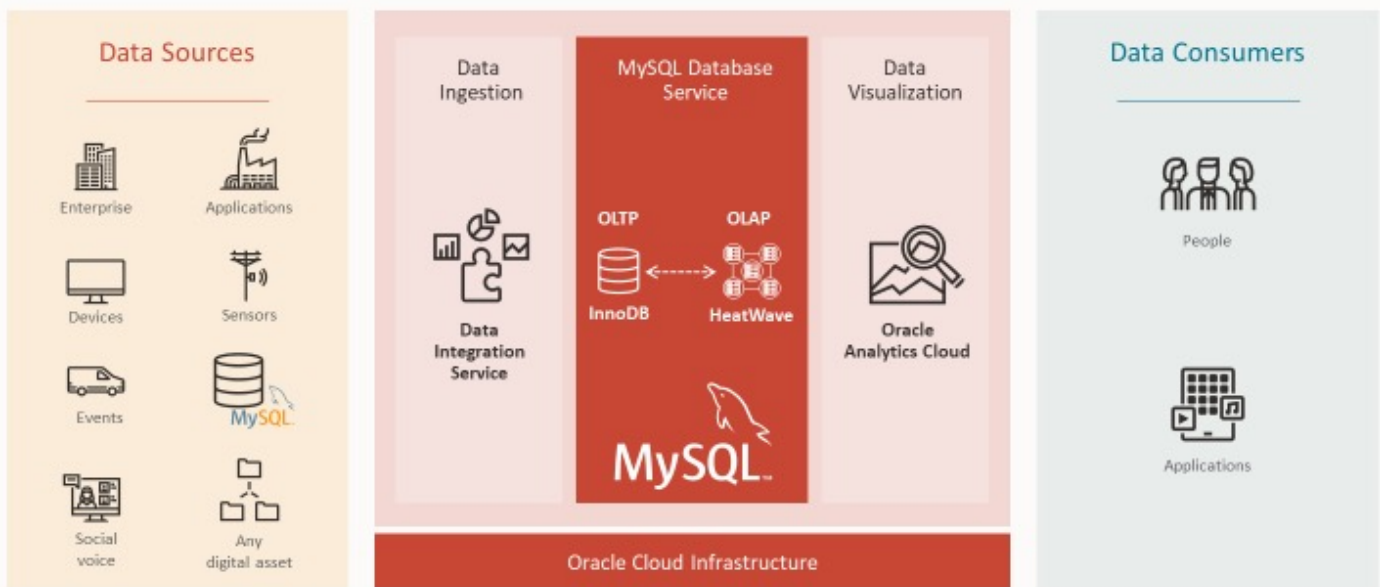
For enterprises to leverage the full efficiencies and synergies of a cloud vendor, it is key that all offerings are well integrated and operate with each other as first-class citizens. That is not always the case for all offerings, especially if there are alternative offerings, where vendors like to “put the thumb on the scale,” typically for the more profitable offering. Not so with Oracle, which gives Oracle MySQL the same “right to play” as its flagship Oracle Database. This is extremely important for Oracle MySQL customers to see so that their confidence in the support and future of Oracle MySQL is reassured.

Specifically, Oracle MySQL Database Service is integrated on four levels (see Figure 10):

Figure 10. How Oracle MySQL Database Service Is Integrated With Other OCI Services

MySQL Database Service is integrated with other Oracle Services

End-to-end Integration from data ingestion to data visualization



Source: Oracle

- **Availability for all data consumers.** Oracle MySQL Database Service is available to all data consumers, whether those consumers are people or applications.
- **Access to all OCI data sources.** All OCI-supported data sources can be processed by Oracle MySQL Database Service: data from enterprises, applications, devices, sensors, events, social, voice, and any digital asset, as well as other MySQL databases.
- **Full support of all OCI data-ingestion capabilities.** Oracle MySQL Database Service is fully compatible with Oracle Data Integrator on OCI, allowing users to leverage Oracle MySQL Database Service the same way as any other Oracle product integrated with Oracle Data Integrator.
- **Full support of all OCI data visualization capabilities.** All visualization from a programming and tooling side is supported as well as on the SaaS side, as is integration with Oracle Analytics Cloud.

ANALYSIS AND OBSERVATIONS

Constellation sees the following strengths and weaknesses for Oracle's MySQL HeatWave offering (see Figure 11):

Strengths

- **Excellent query performance to accelerate insight to action.** Enterprises need to accelerate and become more agile, and that starts with decision management. Insight to action has long eluded CxOs, and MySQL Database Service with HeatWave makes it tangible, enabled by its backward compatibility, high performance, low cost, and with subsecond transactional changes reflected in analytical insights.
- **Ready with zero code changes.** The backward compatibility with no need to change existing OLTP code comprises an elegant architectural approach and enables an immediate uptake of the MySQL Database Service with HeatWave capabilities on top of MySQL transactional data.

- **Fully integrated with other Oracle services.** With Oracle MySQL Database Service being fully integrated with other Oracle Cloud services on OCI, the ecosystem synergies fully come into play. The full integration also eases potential concerns from CxOs toward the future investment in and viability of Oracle MySQL Database Service.
- **Another offering with Oracle's TCO DNA.** Oracle's corporate DNA is all about TCO—providing better performance at the same, often even lower, cost. Oracle MySQL Database Services follows in those footsteps, with leading category performance while offering lower costs.

Weaknesses

- **A new offering.** As with every new product or offering, a healthy degree of technical caution and procedural skepticism is appropriate. This is also valid for Oracle MySQL Database Service with HeatWave, which new customers should test thoroughly for both scale and functionality.
- **Only available in Oracle Cloud Infrastructure.** This may well be only a starting point at the moment (for instance, Oracle makes it possible for customers to use Oracle Database in Microsoft Azure), but for now MySQL Database Service with HeatWave is available only on OCI. Enterprises want to avoid cloud lock-in and therefore favor offerings that support/are available in multiple clouds.
- **It is built. Will they come?** As with any new offering that requires programmatical uptake, it has to be built first and then enterprises can see what the uptake will be. So, only the future will tell which analytical workloads will be run best and first by HeatWave, and what third-party ISV providers will take up the new service.

Figure 11. MySQL HeatWave Strengths and Weaknesses

| STRENGTHS | WEAKNESSES |
|---|--|
| <ul style="list-style-type: none">• Excellent query performance to accelerate insight to action• Ready with zero code changes• Fully integrated with other Oracle Cloud services• Another offering with Oracle's TCO DNA | <ul style="list-style-type: none">• A new offering• Only available in Oracle Cloud• It is built. Will they come? |

Source: Constellation Research

RECOMMENDATIONS

Constellation has the following recommendations regarding the MySQL Database Service with HeatWave:

- 1. Enable Enterprise Acceleration.** Enterprises need to move faster than ever before, and IT/computing infrastructures cannot continue to be the shackles on agility that they have been in the past. Therefore, CxOs should look at any information technology that allows their enterprise to accelerate. Not having to change code to adopt MySQL Database Service with HeatWave is a major benefit and will help the adopting enterprise to move faster and accelerate.
- 2. Enable and then practice insight to action.** Enterprise decision-making has been hampered by the delay of OLTP data being available in OLAP systems. MySQL Database Service with HeatWave puts an end to this from a technology perspective. With the subsecond inclusion of transactional changes in analytical decisions, CxOs can equip their fellow CxOs and decision-makers with an analytical platform that allows real-time insight-to-action best practices.
- 3. Customers using MySQL or MySQL-compatible databases should evaluate MySQL Database Service with HeatWave sooner rather than later.** The immediate benefits of adopting MySQL Database Service with HeatWave are so compelling that existing Oracle MySQL customers should

immediately adopt the new offering. The benefits clearly outweigh the risks, and putting insight to action in reach for CxOs and practicing Enterprise Acceleration is a survival strategy for enterprises.

- 4. Non-MySQL customers need to do a cost-benefit analysis for a potential switch to Oracle MySQL Database Service.** Oracle has achieved quite an engineering feat with MySQL Database Service with HeatWave, which makes it not only a compelling MySQL database offering but also a potential database to migrate workloads to. Although it is possible to try and copy the Oracle technological approach to bring OLTP and OLAP together, it requires serious talent and support to provide a long-lasting, future-proof, and trusted platform.
- 5. Practice commercial prudence.** As always, CxOs need to practice commercial prudence when it comes to platform decisions. One-time costs, ongoing costs, and lock-in effects are the key areas to consider before making platform decisions. Database platform decisions are no exception to the consideration of commercial prudence in all phases of the buying, adoption, and usage cycle.

RELATED RESEARCH

For the CxO perspective on Enterprise Acceleration, see: Holger Mueller, “Why the C-Suite Must Embrace Enterprise Acceleration,” May 2, 2019. <https://www.constellationr.com/research/why-c-suite-must-embrace-enterprise-acceleration>

For the people-leader perspective on the skills shortage and the need for Enterprise Acceleration, see: Holger Mueller, “Why People Leaders Must Embrace Enterprise Acceleration,” Constellation Research, July 3, 2018. <https://www.constellationr.com/research/why-people-leaders-must-embrace-enterprise-acceleration>

For the Market Overview on databases, see: Holger Mueller, “Next-Gen Database Market Overview,” Constellation Research, November 13, 2017. <https://www.constellationr.com/research/next-gen-database-market-overview>

For the importance of Infinite Computing, see: Holger Mueller, “Infinite Platforms Power Enterprise Acceleration,” Constellation Research, October 11, 2019. <https://www.constellationr.com/research/infinite-platforms-power-enterprise-acceleration>

For more details on Infinite Computing, see: Holger Mueller, “The Era of Infinite Computing Triggers Next-Generation Applications,” Constellation Research, June 1, 2018. <https://www.constellationr.com/research/era-infinite-computing-triggers-next-generation-applications>

For more on Business Acceleration Clouds (BACs), see: Holger Mueller, “Enterprise Acceleration Creates Imperative for Business Acceleration Clouds,” Constellation Research, February 7, 2020. <https://www.constellationr.com/research/enterprise-acceleration-creates-imperative-business-acceleration-clouds>

For the Market Overview of next-gen compute platforms, see: Holger Mueller, “Next-Gen Computing: The Enterprise Computing Model for the 2020s,” Constellation Research, September 14, 2018. <https://www.constellationr.com/research/next-gen-computing-enterprise-computing-model-2020s>

For a Constellation ShortList™ of next-generation compute platforms, see: Holger Mueller, “Constellation ShortList Next-Generation Computing Platforms,” Constellation Research, February 12, 2020. <https://www.constellationr.com/research/constellation-shortlist-next-generation-computing-platforms-1>

For an Offering Overview of Oracle Exadata, see: Holger Mueller, “Oracle Exadata Powers Next-Gen Computing,” Constellation Research, June 12, 2019. <https://www.constellationr.com/research/oracle-exadata-powers-next-gen-computing>

RELATED RESEARCH CONTINUED

For an Offering Overview of Oracle Exadata X8M, see: Holger Mueller, “Oracle Exadata X8M: The Fastest Oracle Database Platform,” Constellation Research, September 25, 2019. <https://www.constellationr.com/research/oracle-exadata-x8m-fastest-oracle-database-platform>

For an Offering Overview of Oracle Cloud@Customer, see: Holger Mueller, “Oracle Gen 2 Exadata Cloud at Customer Powers Next-Gen Computing,” Constellation Research, October 7, 2019. <https://www.constellationr.com/research/oracle-gen-2-exadata-cloud-customer-powers-next-gen-computing>

For next-generation databases, see: Holger Mueller, “Constellation ShortList Next-Gen Databases: RDBMS for On-Premises,” Constellation Research, February 12, 2020. <https://www.constellationr.com/research/constellation-shortlist-next-gen-databases-rdbms-premises-3>

For a Constellation ShortList of low-code tools and platforms, see: Dion Hinchcliffe and Holger Mueller, “Constellation ShortList Enterprise Low-Code Tools and Platforms,” Constellation Research, February 19, 2020. <https://www.constellationr.com/research/constellation-shortlist-enterprise-low-code-tools-and-platforms-4>

For a Constellation ShortList on IaaS vendors, see: Holger Mueller, “Constellation ShortList Global IaaS for Next-Gen Applications,” Constellation Research, February 19, 2020. <https://www.constellationr.com/research/constellation-shortlist-global-iaas-next-gen-applications-4>

For the success of open source, see: Holger Mueller, “How Open Source Won and Will Keep Winning,” August 20, 2019. <https://www.constellationr.com/research/how-open-source-won-and-will-keep-winning>

On the future of code automation, see: Holger Mueller, “The Future of DevOps Is Automated Software Operations,” February 12, 2020. <https://www.constellationr.com/research/constellation-shortlist-devops-0>

For more best-practice considerations for PaaS offerings, see: Holger Mueller, “As PaaS Turns Strategic, So Do Implementation Considerations,” Constellation Research, May 9, 2018. <https://www.constellationr.com/research/paas-turns-strategic-so-do-implementation-considerations>

For more on next-gen applications and PaaS offerings, see: Holger Mueller, “Why Next-Gen Apps Start With a Next-Gen Platform as a Service,” April 5, 2018. <https://www.constellationr.com/research/why-next-gen-apps-start-next-gen-platform-service>

For a Constellation ShortList on PaaS vendors, see: Holger Mueller, “Constellation ShortList PaaS Tool Suites for Next-Gen Apps,” Constellation Research, February 19, 2020. <https://www.constellationr.com/research/constellation-shortlist-paas-tool-suites-next-gen-apps-3>

RELATED RESEARCH CONTINUED

Also: Holger Mueller, "Constellation ShortList PaaS Suites for Next-Gen Apps," February 27, 2018. <https://www.constellationr.com/research/constellation-shortlist-paas-suites-next-gen-apps-3>

For additional IaaS and PaaS selection criteria, see: R "Ray" Wang and Holger Mueller, "Key Questions for Every Public Cloud IaaS/PaaS Decision Matrix," Constellation Research, January 24, 2018. <https://www.constellationr.com/research/key-questions-every-public-cloud-iaaspaas-decision-matrix>

For an example of digital transformation, see: Holger Mueller, "Lufthansa Digitally Transforms the Workplace for Flight Managers," Constellation Research, February 27, 2018. <https://www.constellationr.com/research/lufthansa-digitally-transforms-workplace-flight-managers>

For an example of digital transformation, see: Holger Mueller, "How Oracle Graal Supercharged Twitter's Microservices Platform," Constellation Research, December 11, 2019. <https://www.constellationr.com/research/how-oracle-graalvm-supercharged-twitter-s-microservices-platform>

ENDNOTES

¹ For more on the history of MySQL, see: MySQL, Wikipedia: <https://en.wikipedia.org/wiki/MySQL>

² Constellation refers to “Enterprise Acceleration” as the need for enterprises to move faster and become more agile. For more, see: Holger Mueller, “Why the C-Suite Must Embrace Enterprise Acceleration,” May 2, 2019. <https://www.constellationr.com/research/why-c-suite-must-embrace-enterprise-acceleration>

ANALYST BIO

Holger Mueller

Vice President and Principal Analyst

Holger Mueller is vice president and principal analyst at Constellation Research, providing guidance for the fundamental enablers of the cloud, IaaS, and PaaS, with forays up the tech stack into big data, analytics, and SaaS. Mueller provides strategy and counsel to key clients, including chief information officers (CIOs), chief technology officers (CTOs), chief product officers (CPOs), investment analysts, venture capitalists, sell-side firms, and technology buyers.

Prior to joining Constellation Research, Mueller was VP of products for NorthgateArinso, a KKR company. He led the transformation of products to the cloud and laid the foundation for new business-process-as-a-service (BPaaS) capabilities. Previously, he was the chief application architect with SAP and was also VP of products for FICO. Before that, he worked for Oracle in various management functions—on both the application development (CRM, Fusion) and business development sides. Mueller started his career with Kiefer & Veitinger, which he helped grow from a startup to Europe's largest CRM vendor from 1995 onward. Mueller has a Diplom-Kaufmann degree from the University of Mannheim, with a focus on information science, marketing, international management, and chemical technology. A native European, Mueller speaks six languages.

 [@holgermu](https://twitter.com/holgermu)  constellationr.com/users/holger-mueller  linkedin.com/in/holgermueller

ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley–based research and advisory firm that helps organizations navigate the challenges of digital disruption through business model transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered, and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors, and vendor clients. Our mission is to identify, validate, and share insights with our clients.

Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management, and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.



www.ConstellationR.com



[@ConstellationR](https://twitter.com/ConstellationR)



info@ConstellationR.com



sales@ConstellationR.com

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, emailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research Inc. is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a committee of publishers and associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco Bay Area | Boston | Colorado Springs | Ft. Lauderdale | Los Angeles | New York Metro
Northern Virginia | Portland | Pune | Sacramento | San Diego | Sydney | Toronto | Washington, D.C.