

October 6, 2023



OFFERING OVERVIEW

# Oracle MySQL HeatWave Grows, Adds Lakehouse Support

One database optimized for analytics, machine learning,  
and OLTP



**Holger Mueller**  
VP & PRINCIPAL ANALYST



*Produced exclusively for Constellation Research clients*

# TABLE OF CONTENTS

- Executive Summary .....3
- About Oracle MySQL HeatWave .....4
- Functional Capabilities..... 10
- Analysis and Observations ..... 18
- Recommendations ..... 21
- Related Research ..... 23
- Related Research Continued..... 24
- Related ShortLists..... 25
- Endnotes ..... 26
- Analyst Bio..... 27
- About Constellation Research ..... 28



# EXECUTIVE SUMMARY

This Offering Overview examines the new capabilities of Oracle MySQL HeatWave, which added lakehouse support in summer 2023. The report describes the underlying market trends, introduces the vendor, and presents key differentiators for Oracle's offering. It continues with an analysis of the vendor's strengths and weaknesses and concludes with a set of tangible and actionable recommendations for CxOs.<sup>1</sup>

## BUSINESS THEMES



New C-Suite



Technology  
Optimization



Data to Decisions

# ABOUT ORACLE MYSQL HEATWAVE

## Overview

MySQL became part of Oracle via the Sun Microsystems acquisition in 2009.<sup>2</sup> Oracle has been a good steward of MySQL, supporting and enhancing the MySQL platform since the acquisition and adding new features and capabilities, such as a high-availability (HA) version. Oracle started its MySQL innovation super cycle in December 2020 with the addition of HeatWave, a cloud-native in-memory query accelerator designed to speed online analytical processing (OLAP) to deliver real-time analytics and other complex queries within the MySQL database as a managed service in Oracle Cloud Infrastructure (OCI).

Oracle has continued its investments and innovations for MySQL, adding machine learning (ML)-based automation to HeatWave with a component called MySQL Autopilot in August 2021. Autopilot brings management and performance automation capabilities without any additional charge, reducing the need for many time-consuming manual tasks usually performed by DBAs. And in spring 2022, Oracle expanded the HeatWave service further, adding HeatWave AutoML to the already available online transaction processing (OLTP) and OLAP capabilities inside MySQL HeatWave and automating the ML lifecycle to accelerate model development and make it less of a data scientist domain. Oracle also announced more innovation with real-time elasticity and doubled the amount of data processed in a MySQL HeatWave node to reduce costs, added support for database pause and resume, and provided new benchmark results.

In fall 2022, Oracle announced the availability of a native implementation of HeatWave on AWS as well as new automation capabilities in Autopilot, and enterprise-class security features with significant price/performance numbers. Finally, Oracle just added lakehouse support to HeatWave as well as lakehouse-specific innovations to Autopilot. This makes MySQL HeatWave the only database service that can perform OLTP and real-time OLAP, and machine learning, as well as offer big data characteristics with lakehouse support that can handle up to a 500TB data lake. Indeed, leveraging

AMD EPYC™ processors, HeatWave Lakehouse on OCI demonstrated orders of magnitude faster performance than competitive offerings on a 500TB TPC-H benchmark.

## 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 that database choices matter. Why? Switching databases is both hard and expensive, and doing so often poses a risk to business continuity and security. 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

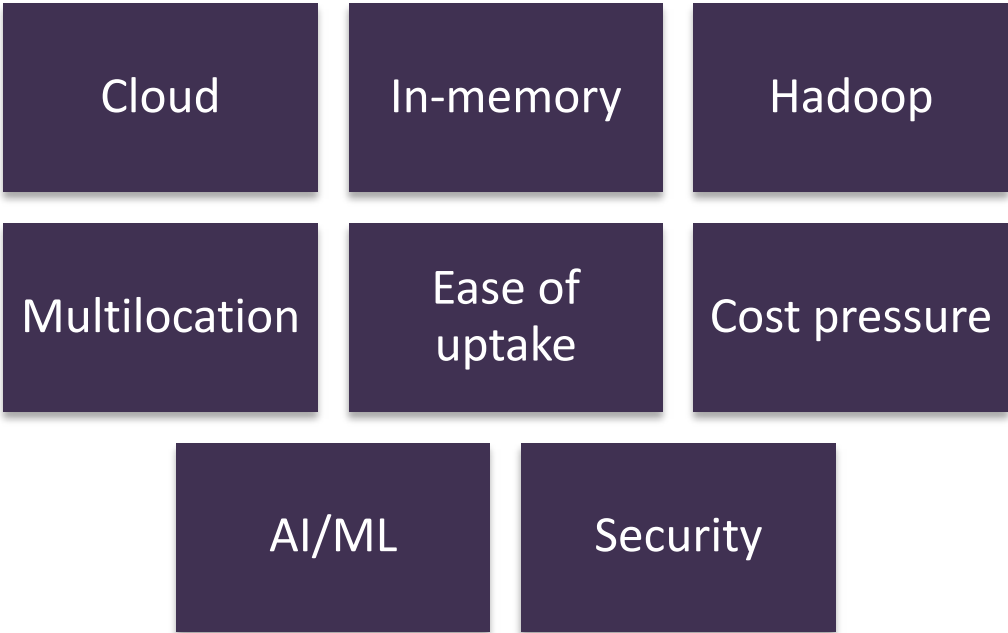
Constellation has identified seven key market trends for next-generation databases (see Figure 1):

- 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, hosting databases in the cloud 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 in Constellation's associated Next-Gen Database Market

Overview<sup>3</sup>—except 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.

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 (relatively expensive) memory, but some also offer in-memory options. Consequently, how memory is managed, accessed, and consumed is where the vendors differ. Hadoop-based vendors use Apache Spark and Hive. Traditional relational database management system (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 always place the entire system in memory. Meanwhile, cloud-based vendors, given the novelty

**Figure 1. Seven Market Trends Defining a Next-Gen Computing Platform**



Source: Constellation Research

of their offering, have not shared their in-memory philosophy. Constellation expects the uptake of in-memory options to be like 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 could store all their electronic information in one place, without knowing what they want to analyze and while remaining commercially viable. In addition, a data storage and retrieval architecture can bridge data silos. This approach to big data is now referred to as a data lake, now mostly stored in the cloud.
- 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, through either 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 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 traditionally expensive

database licenses. All traditional database vendors can already see the pressure on their license prices. Constellation expects this trend to continue.

- 7. AI and ML change database capabilities dramatically.** The rise of AI/ML has dramatically changed the way databases are operated and used. On the operation side, databases now run AI/ML processes on their operations processes, increasing uptime and reliability and lowering operating costs. For developers, databases make AI/ML constructs available that allow to build AI/ML-enabled next-generation applications in an easier, more reliable, and faster way. And finally, AI/ML changes the user experience as users now have access to AI/ML-based automation as well as easy-to-use constructs to apply them in their daily database usage.
- 8. 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 workloads 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 and OLAP Versus Big Data Debate Ends With Lakehouse Technology

For the longest time, enterprise data processing was based on OLTP. Then came OLAP, and about a dozen years ago came big data. The qualities of OLTP/OLAP versus big data were fundamentally different—from response time and access to cost of the solution.

In the process, big data got a bad reputation, because its solutions were hard to manage and difficult to keep going due to the volatility of the underlying open source. So, the industry turned to reinventing the big data offerings by renaming them. It looks as if the industry has settled now on “lakehouse”—a combination of data warehouse (from the OLAP world) and data lake (from the big




data domain). Today, enterprises accept the necessity of operating lakehouses to provide them the insights on all of their data. The main reason is that the largest part of enterprise data is not in the OLTP and OLAP systems, so enterprises need to access all of their data—including logs, documents, and communication files (see Figure 2).

To a certain point, Oracle MySQL HeatWave has followed the progression of the industry, starting with OLTP support and then adding OLAP support, then differentiated with database machine learning and more. It seems only logical that the HeatWave team has now added support for lakehouses. That matters, because more than 80% of data is stored in file systems and 99.5% of the data enterprises collect and store remains unused. All of this data is outside of the solution domain of OLTP and traditional OLAP tools. With MySQL HeatWave’s lakehouse support, the complete set of enterprise and external data becomes accessible and available for insights—with industry-leading performance and cost.

## Figure 2. The Data Diaspora in the Enterprise

### Massive amount of data stored outside databases

- Databases are systems of record
- Files are repository for other types of data (e.g., IoT, web content, log files)
- **99.5%** of collected data remains unused



1 Copyright © 2023 Oracle and/or its affiliates. | Oracle Confidential – Includes non-disclosure information that must be covered by a current NDA agreement.



Source: Oracle

MySQL HeatWave's unrivaled performance is a result of its scale-out architecture that leverages AMD EPYC™ processors powering massive parallelism to provision the cluster, load data, and process queries with up to 512 nodes. MySQL HeatWave takes advantage of the large L3 cache size and DRAM memory bandwidth available on AMD EPYC™ to deliver unparalleled database and in-memory analytics performance. In addition, enhancements to MySQL Autopilot automate metadata creation for object files and dynamically adapt to the performance of the underlying object store to provide the best performance in any OCI region.

## FUNCTIONAL CAPABILITIES

### A Universal Database Becomes a Universal Data Processing Platform

One of the first innovations Oracle delivered with MySQL HeatWave was the combination of OLTP and OLAP in a single database service. With that key combination, the offering moved to being a universal database,<sup>4</sup> offering more than a single-function database can provide, but it did not stop there.

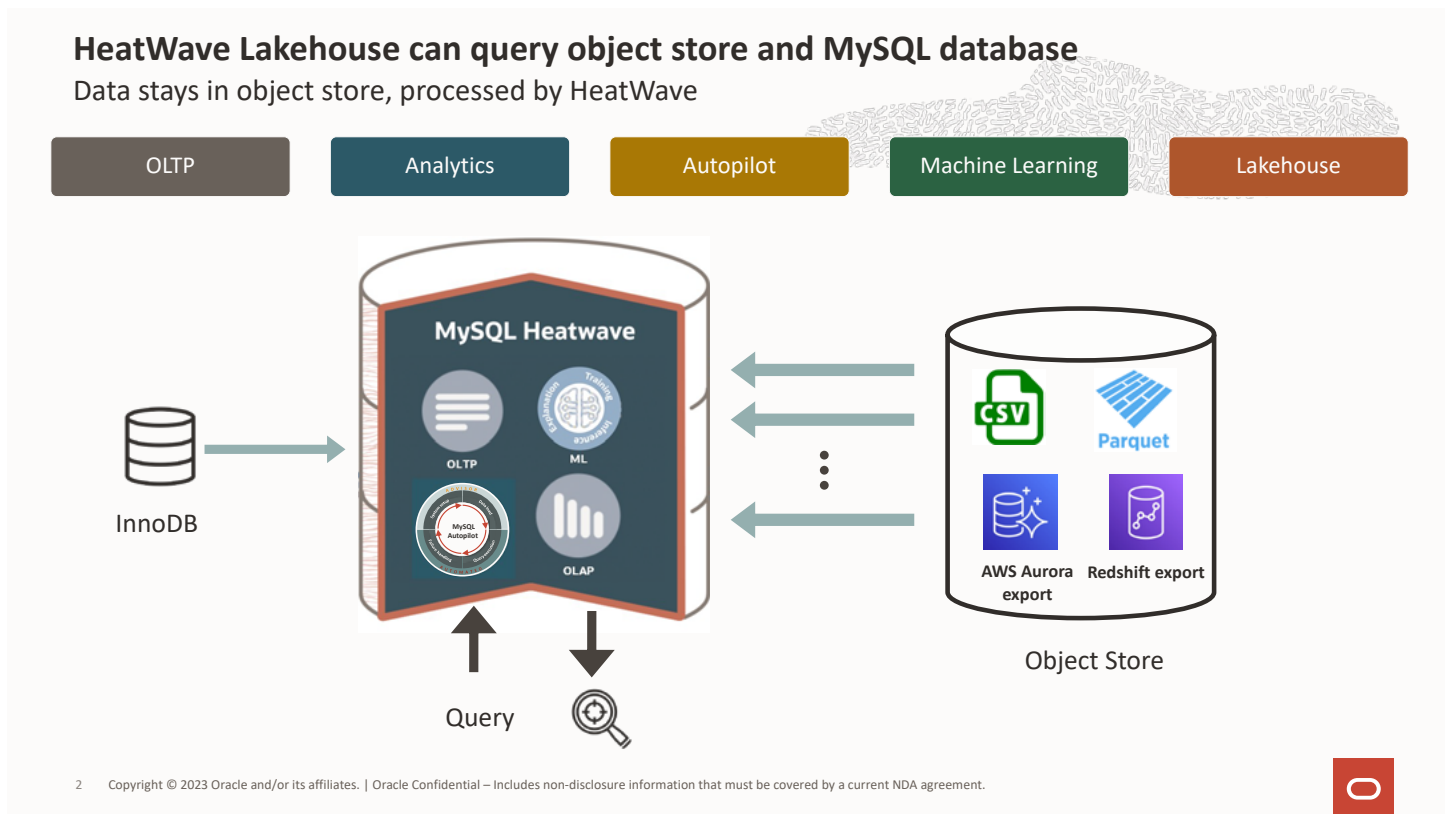
Specifically, MySQL HeatWave now offers the following:

- **One database for OLTP and OLAP.** One of the key value drivers for MySQL HeatWave is the ability to run OLTP as well as OLAP workloads—all in a single database service, eliminating extract, transform, and load (ETL) processes and enabling real-time analytics.<sup>5</sup>
- **Built-in support for machine learning with HeatWave AutoML.** MySQL HeatWave ended the separation of database data and data used to model ML applications, eliminating ETL and enabling models based on real-time data (sound familiar?). ML models can be built on HeatWave nodes via HeatWave AutoML, which fully automates the training process, inference, and explanation of the ML models.<sup>6</sup>
- **Self-driving capabilities with MySQL Autopilot.** Manual database administration is a best practice of the past: MySQL Autopilot runs key database administrative functions for customers, thus

freeing up the DBA's time for more high-value activities.

- **Multicloud support.** These days customers are using multiple clouds, taking the best services of each. HeatWave is available natively on OCI, AWS, on-premises with OCI Dedicated Region, and to Azure users via Oracle Database Service for Azure. Needless to say, all of these cloud offerings are fully managed by Oracle and have Autopilot capabilities. It is also important to note that the Oracle support teams have direct access to the MySQL product development team. This gives peace of mind to CxOs using HeatWave on any cloud.
- **One database for relational, analytical, and lakehouse needs.** With Oracle's latest innovation as of summer 2023, HeatWave adds support for object stores, the platform of lakehouses (see Figure 3).

**Figure 3. How HeatWave Works With an Object Store**



Source: Oracle

- **Support for MySQL-compatible and non-MySQL workloads.** With HeatWave Lakehouse, the processing is done in the HeatWave engine and the data is not copied into the MySQL database. Customers migrating from non-MySQL databases can upload their data to the object store and leverage all of the functional, performance, and scalability benefits of HeatWave.

The HeatWave team decided to leave data inside of the object store, avoiding redundant data storage. For MySQL HeatWave to query the data, the database creates a representation of the lakehouse data, effectively mapping any data schema and making it accessible to HeatWave's query plan. Effectively, HeatWave does not “know” anymore where the data is that it queries, because the lakehouse representation inside of HeatWave is seamless. The result is something remarkable for big data-class solutions: identical query time for querying data coming from the database or the object store.

The result of all of this is an extremely attractive HeatWave offering. It now not only is a universal database but also is a universal data processing platform that supports all critical enterprise data processing needs in a very effective way.

## MySQL Autopilot Is HeatWave's Secret Superpower

The ability to use AI/ML-based automation in software offerings has quickly turned out to be a key success factor for automating critical processes. MySQL HeatWave is a great showcase for using AI/ML to automate important database tasks, adding continuous capabilities to its Autopilot capabilities.

Consequently, it comes as no surprise that the HeatWave team not only enhanced existing capabilities, but also added new Autopilot capabilities leveraging AI/ML heavily to automate the key tasks associated with lakehouses, as follows (see Figure 4):

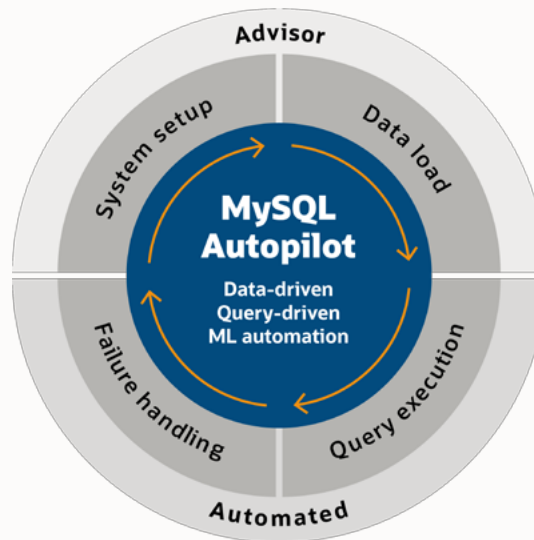
- **Auto provisioning makes system management easy.** Setting up and operating a database can be tedious and time-consuming—and risky if an error is made. HeatWave makes that a challenge of the past with its auto-provisioning capabilities in Autopilot that predict the optimal size of the cluster to meet performance requirements.

- **Auto schema inference makes DBAs' life simpler.** To process data, databases need to understand it, typically by knowing the schema of the underlying data. Creating schema manually is time-consuming and can include errors that could have performance-degrading downsides. HeatWave's auto schema inference makes this process a non-challenge—and automatically determines the data type, precision, and length of the attributes in the files.
- **Adaptive data sampling improves predictions.** MySQL Autopilot intelligently samples files to derive information about the data that is needed for making Autopilot predictions. With this capability, MySQL Autopilot can scan and make schema predictions on a 400TB file in under one minute.
- **Auto parallel loading ensures high performance.** When it comes to loading data, parallelism is highly desirable, but also hard to set up, manage, and operate. Autopilot automates the tasks, making it easy to load data as fast as possible and as fast as needed.

**Figure 4. MySQL Autopilot Capabilities for Lakehouse**

### Autopilot for HeatWave Lakehouse

- Auto provisioning
- Auto schema inference
- Adaptive data sampling
- Auto parallel loading
- Adaptive data flow
- Auto query plan improvement



Source: Oracle

- **Adaptive data flow enables optimal performance.** MySQL Autopilot coordinates bandwidth utilization to the object store across a large cluster of nodes, dynamically adapting to the performance of the underlying object store, resulting in optimal performance and availability in any OCI region.
- **Auto query plan increases performance.** With Autopilot constantly improving the query plans—now including the data in the lakehouse—high performance is ensured for any data queries an enterprise may need HeatWave to process.

In summary, the capabilities of Autopilot are essential for an enterprise to succeed with HeatWave. It continues the long tradition of Oracle using ML to make HeatWave more robust, resilient, high-performing, and easier to use.

## How HeatWave Delivers Unbeaten Lakehouse Query Performance

For MySQL HeatWave to be able to efficiently process data in object store, it needs to gather statistics and information about the data. The most critical process here is the initial load/ingestion, because later loads are smaller when it comes to data volume.

HeatWave achieves impressive data-loading performance thanks to its ability to scale out the data-loading process. By separating the data into “chunks,” HeatWave is able to dynamically process resources to the workload. In the process, it collects statistics and aggregates them, transforms and compresses data, stores the optimized format in the object store, and moderates the workflow demands based on available bandwidth.

Databases always struggle when it comes to resource contention caused when too many queries are coming from too many users. To make things worse, sometimes those queries have dependencies and hold each other up, significantly reducing overall database performance.

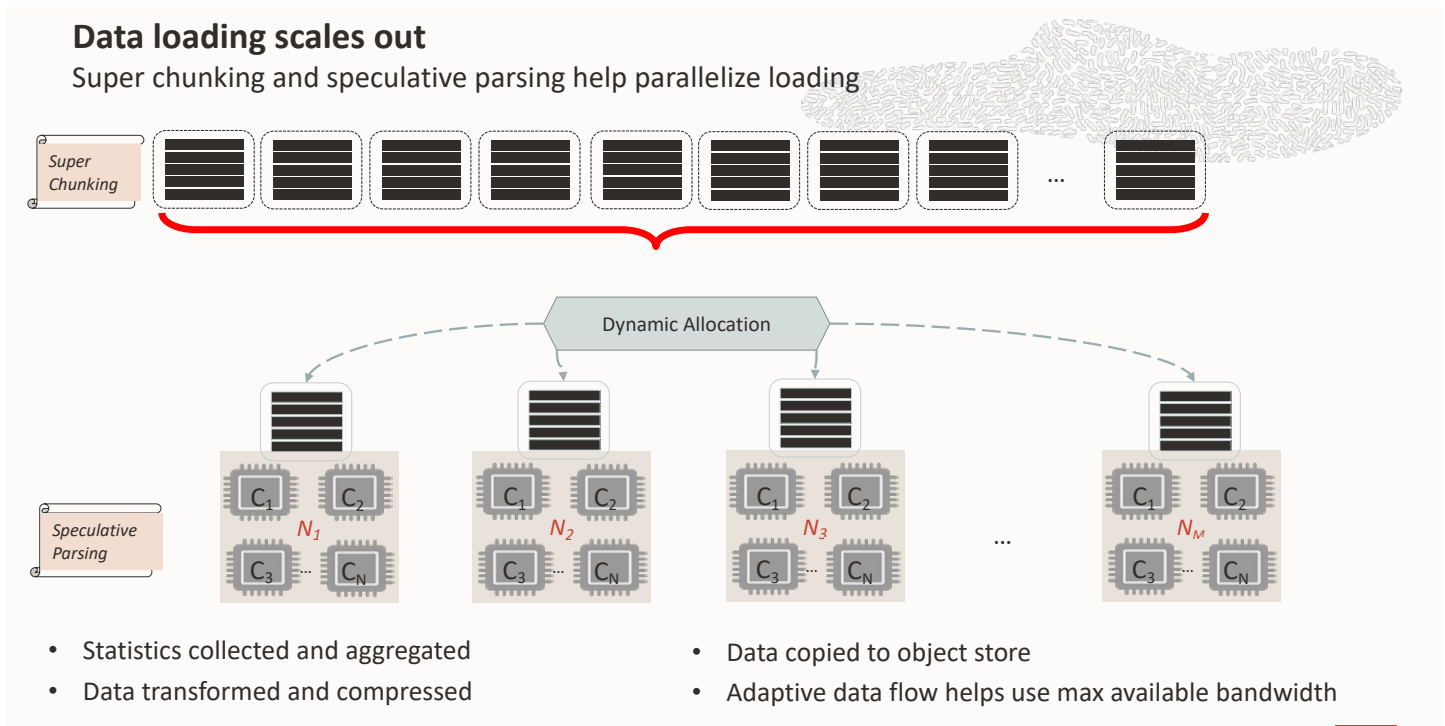
The HeatWave team used ML to tackle the problem, making Autopilot aware of the transactional workload running in the system. With that awareness, HeatWave can unblock resource contentions

that are an issue for traditional databases' work scheduling engines, effectively aligning transactions for overall optimized performance (see Figure 5).

The result of this architecture is industry-leading load performance that beats competitive architectures and offerings by factors of two to nine. And this higher performance helps enterprises not only to get to their insights faster and earlier but also to keep related costs down. Faster means cheaper in the cloud era.

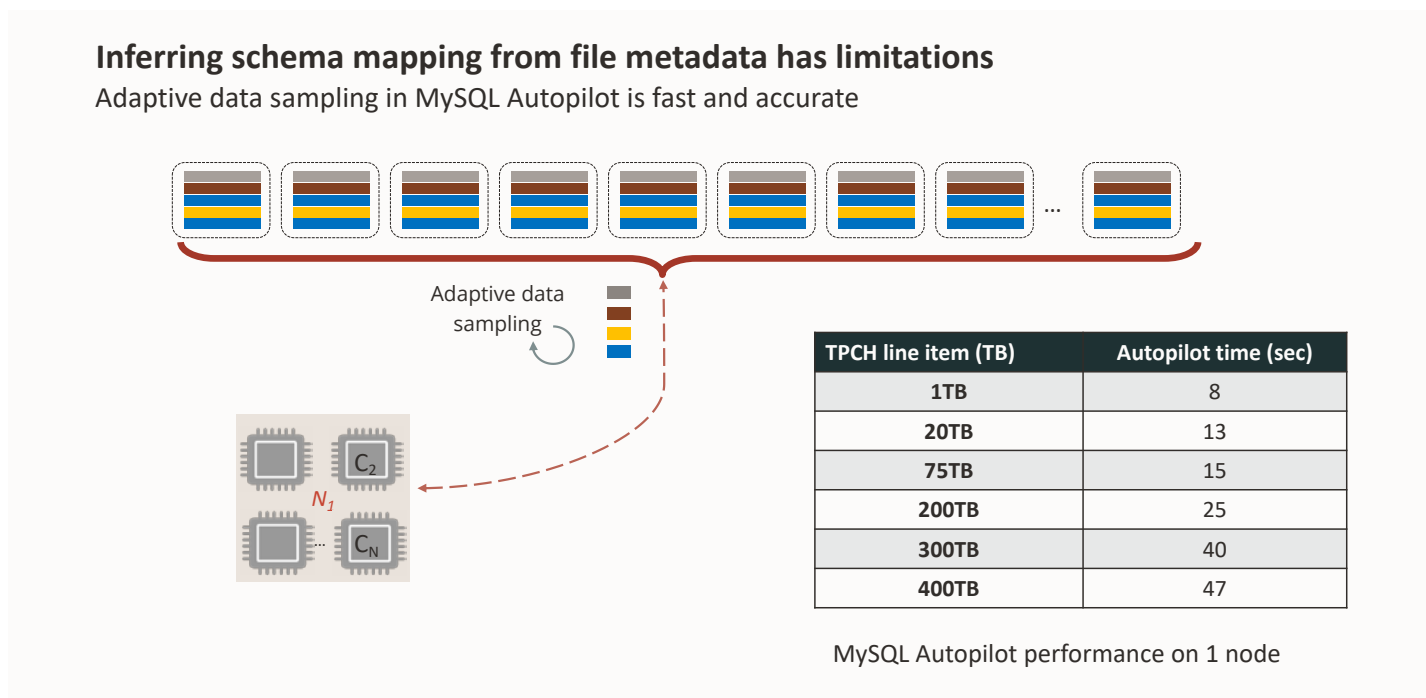
As if this were not enough, HeatWave also overcomes one of the main drawbacks of automatic schema mapping, which is inaccurate mapping (see Figure 6). HeatWave overcomes this by intelligently sampling the data (not just reading from the header file), which results in accurate mapping for the data to the schema information. The intelligent sampling results in great performance—a 400TB file can be scanned in under a minute.

**Figure 5. How HeatWave Addresses Data Scale-Out**



Source: Oracle

## Figure 6. The Limits of Schema Mapping



Source: Oracle

Overall, HeatWave has an architecture that is uniquely suited to manage ingestion, schema mapping, and sampling of an object store to able it to query it efficiently with high performance, and cost-consciously.

## A Simplification of the Enterprise Data Landscape With Amazing Performance

Enterprises have always struggled to unify their data landscape. The result has been slower reaction time, slower and sometimes bad decisions, and lack of visibility despite having (and paying for the storage of) the data—in short, all detractors from what matters to CxOs: Enterprise Acceleration<sup>7</sup> (see Figure 7).

MySQL HeatWave puts an end to this, bringing together all enterprise data that is stored, regardless of whether it is in OLTP, OLAP, data warehouse, or an object store or lakehouse. And it does that at a remarkably low cost and with better performance than the usual competition. The successful use of AI/ML inside the database has been expanded to efficiently manage the new object storage. And that



## Figure 7. HeatWave's KPIs Running a Lakehouse

### Simplify data management with HeatWave Lakehouse

For all data sizes, operational system or data warehouse

#### MySQL Autopilot:

1TB in 8 sec, 400TB file in under 1 minute

#### Provision HeatWave:

1 node in 12 minutes, 512 nodes in 16 minutes

#### Load:

1TB in 1.1 hours, 500TB in 4.4 hours

#### Average Query time:

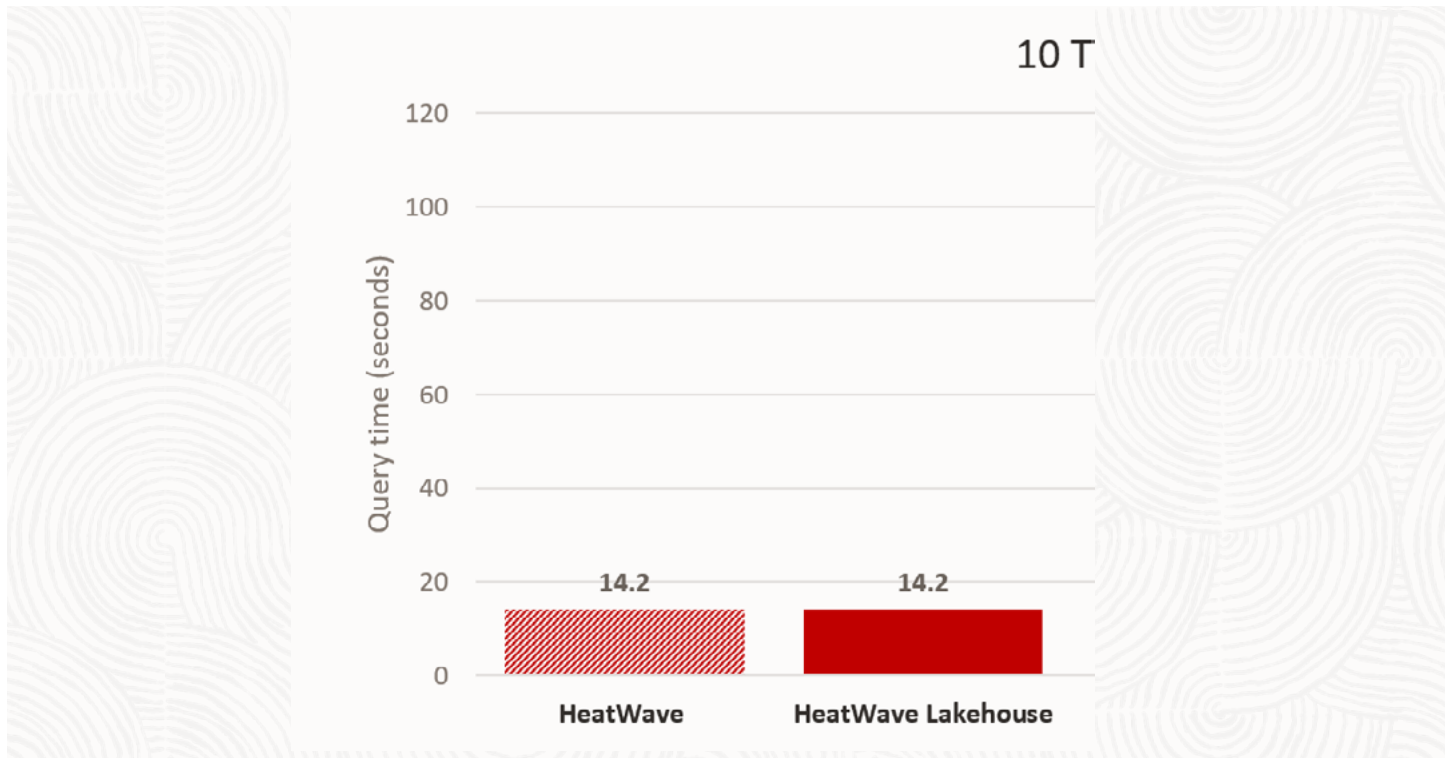
6 sec for 1TB workload, 47 sec for 500TB workload

Source: Oracle

matters to CxOs, because it does not require a small army of highly skilled, highly compensated, and hard-to-retain data professionals. Autopilot automates all critical aspects of HeatWave, to the point that even a tech-savvy business user could add/load/query lakehouse data in HeatWave.

The most remarkable quality of Oracle's solution is the parity in performance of querying data inside the database and in object store. It is remarkable because until now big data-style technologies—such as the lakehouse—were always slower than querying data inside a database. The first Hadoop clusters would come back in a few minutes. Today, response times tend to be below 30 seconds—a massive improvement but still a far cry from OLTP performance. Now, HeatWave manages to match its OLTP performance to the identical level for the lakehouse. This is an industry first, and a testament to the skills and tech savviness of the HeatWave team—as well as to the efficient implementation of Oracle Cloud (e.g., network speeds play a role). And with identical performance also comes identical cost, which is equally remarkable (see Figure 8).

**Figure 8. HeatWave’s Identical Performance for OLTP and Lakehouse (10TB TBC-H Query Performance)**



Source: Oracle

The result of this is that, for the first time ever, big data with its large data volume is being analyzed at the same speed as OLTP data. This opens completely new horizons in data discovery, insights, and automation. And with that, HeatWave becomes a data platform that CxOs cannot afford to ignore—and, beyond that, need to evaluate—because they can’t have their enterprise miss out on a platform that enables breakthroughs in enterprise insights.

## ANALYSIS AND OBSERVATIONS

Constellation sees the following strengths and weaknesses for HeatWave (see Table 1):

### Strengths

- **A cloud-native data platform.** MySQL HeatWave is much more in 2023 than the MySQL-compatible database it was at its inception only a few years back. Today, with the addition of

lakehouse support, it has become the cloud-native data platform for all data processing needs of an enterprise.

- **One single platform for all data and analytics.** MySQL HeatWave ends the current best practice to separate OLTP/OLAP from lakehouses, instead unifying both worlds—remarkably, with identical query performance and TCO—for unparalleled qualitative insights based on enterprise data.
- **Support for non-MySQL workloads.** With HeatWave Lakehouse, all of the processing is done in the HeatWave engine. As a result, the system can now support processing non-MySQL workloads that customers can load in the object store.
- **Unparalleled automation with MySQL Autopilot.** Autopilot automates, operates, and manages HeatWave to a level of automation not seen in traditional and cloud databases. Autopilot allows enterprises to reduce complexity; eliminates the need for expensive specialists; and manages HeatWave via scalable automation, overcoming the reliance on human operators.
- **It is another testament to Oracle's TCO-focused organizational DNA.** Oracle's corporate DNA is to offer better products at lower TCO. HeatWave is another great showcase documenting that Oracle's corporate DNA is alive and well. HeatWave offers industry-leading performance and record-low TCO.

## Weaknesses

- **New offerings need time to prove themselves.** As with every new product or offering, a healthy degree of technical caution and procedural skepticism is appropriate. This is also valid for HeatWave Lakehouse, which new customers should test thoroughly for both scale and functionality. HeatWave has rapidly matured as a service rich in functionality and is one of the fastest-growing services on OCI.

- **Oracle needs to keep earning—and keeping—the trust of customers.** Oracle is a well-respected leader for transactional databases, cloud applications, and a development/technology stack. When it comes to cloud infrastructure and new offerings in the cloud, the company needs to earn the trust of customers—and keep earning it. Both OCI and HeatWave are on a good trajectory in this regard, but they need to retain their momentum.
- **It is built—time will tell which workloads have the best uptake.** As with any new offering that requires programmatic uptake, it must be built first, and then enterprises can see what the uptake will be. Only the future will tell which analytical and ML workloads will be run best and first by HeatWave in combination with a lakehouse. Independent software vendors (ISVs) will take up the new service and will also be a key factor for future enterprise adoption. Given MySQL HeatWave’s attractive TCO, it is highly likely that workloads will gravitate toward it.

**Table 1. HeatWave Strengths and Weaknesses**

STRENGTHS	WEAKNESSES
<ul style="list-style-type: none"> <li>- A cloud-native data platform</li> <li>- One single platform for all data and analytics</li> <li>- Support for non-MySQL workloads</li> <li>- Unparalleled automation with MySQL Autopilot</li> <li>- Another testament to Oracle’s TCO-focused organizational DNA</li> </ul>	<ul style="list-style-type: none"> <li>- As a new offering, HeatWave will need time to prove itself</li> <li>- Oracle needs to keep earning—and keeping—the trust of customers</li> <li>- It is built—time will tell which workloads have the best uptake.</li> </ul>

**Source:** Constellation Research

# RECOMMENDATIONS

Constellation offers the following recommendations regarding MySQL HeatWave:

- 1. Power Enterprise Acceleration with HeatWave.** Enterprises need to move faster than ever, 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 enables their enterprises to accelerate. Being able to use one single offering, HeatWave, as the platform for all enterprise data needs makes HeatWave a key driver for Enterprise Acceleration.
- 2. Build on an unparalleled data foundation for next-best action.** Enterprise decision-making has been hampered by the delay of OLTP data availability in OLAP systems and the inability to access unstructured data in lakehouses. HeatWave brings all of this together along with unseen levels of high performance as well as low cost, bringing the data foundation of an enterprise into a single offering. This reduces complexity, enables faster insights, and lowers costs—all highly desirable outcomes for any CxO.
- 3. Evaluate HeatWave now.** Customers using MySQL-compatible databases or non-MySQL workloads should assess the immediate benefits of adopting HeatWave in general and of adopting this new offering in particular. Enterprises operating lakehouses now have a chance to bring all of their data together in a single offering. HeatWave's unique architecture and cost advantages make it an appealing offering, even for enterprises that do not operate MySQL.
- 4. Conduct a cost/benefit analysis for a potential switch to HeatWave Lakehouse.** Non-MySQL customers need to run the numbers. Oracle has achieved quite an engineering feat with HeatWave in general and now with its addition of lakehouse processing, which makes it not only an even more compelling database offering but also a potential database to migrate workloads to.
- 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 original launch of Oracle HeatWave, see: Holger Mueller, "There Is a HeatWave for Databases: Oracle MySQL Database Service Merges OLTP and OLAP," Constellation Research, June 3, 2021. <https://www.constellationr.com/research/there-heatwave-databases-oracle-mysql-database-service-merges-oltp-and-olap>

---

For the latest research on Oracle MySQL, see: Holger Mueller, "Oracle HeatWave Brings Machine Learning Inside the Database," Constellation Research, March 29, 2022. <https://www.constellationr.com/research/oracle-mysql-heatwave-brings-machine-learning-inside-database>

---

For the CxO perspective on Enterprise Acceleration, see: Holger Mueller, "Why the C-Suite Must Embrace Enterprise Acceleration," Constellation Research, 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 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 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>

---

## RELATED RESEARCH CONTINUED

---

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>

---

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 the success of open source, see: Holger Mueller, “How Open Source Won and Will Keep Winning,” Constellation Research, 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,” Constellation Research, June 21, 2019. <https://www.constellationr.com/research/future-devops-automated-software-operations>

---

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

---

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,” Constellation Research, April 5, 2018. <https://www.constellationr.com/research/why-next-gen-apps-start-next-gen-platform-service>

---

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>



## RELATED SHORTLISTS™

---

For a Constellation ShortList of next-generation compute platforms, see: Holger Mueller, “Constellation ShortList™ Next-Generation Computing Platforms,” Constellation Research, February 9, 2022. <https://www.constellationr.com/research/constellation-shortlist-next-generation-computing-platforms-1>

---

For a Constellation ShortList of next-generation databases, see: Holger Mueller, “Constellation ShortList™ Next-Gen Databases: RDBMS for On-Premises,” Constellation Research, February 9, 2022. <https://www.constellationr.com/research/constellation-shortlist-next-gen-databases-rdbms-premises-4>

---

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, August 17, 2022. <https://www.constellationr.com/research/constellation-shortlist-enterprise-low-code-tools-and-platforms-5>

---

For a Constellation ShortList on IaaS vendors, see: Holger Mueller, “Constellation ShortList™ Global IaaS for Next-Gen Applications,” Constellation Research, February 16, 2022. <https://www.constellationr.com/research/constellation-shortlist-global-iaas-next-gen-applications-5>

---

For a Constellation ShortList on AI development platforms, see: Holger Mueller, “Constellation ShortList™ AI-Powered Developer Platforms,” Constellation Research, August 17, 2022. <https://www.constellationr.com/research/constellation-shortlist-ai-powered-developer-platforms>

---

For a Constellation ShortList on PaaS vendors, see: Holger Mueller, “Constellation ShortList PaaS Tool Suites for Next-Gen Apps,” Constellation Research, February 9, 2022. <https://www.constellationr.com/research/constellation-shortlist-paas-tool-suites-next-gen-apps-5>

---

Also: Holger Mueller, “Constellation ShortList™ PaaS Suites for Next-Gen Apps,” Constellation Research, February 16, 2022. <https://www.constellationr.com/research/constellation-shortlist-paas-suites-next-gen-apps-4>

## ENDNOTES

---

<sup>1</sup> For more coverage on Oracle's MySQL HeatWave offering, also consider looking at: Holger Mueller, "There Is a HeatWave for Databases: Oracle MySQL Database Service Merges OLTP and OLAP," Constellation Research, June 3, 2021. <https://www.constellationr.com/research/there-heatwave-databases-oracle-mysql-database-service-merges-oltp-and-olap>

Holger Mueller, "Oracle Switches to Autopilot, Turns Up the Heat With MySQL Database Service," Constellation Research, September 16, 2021. <https://www.constellationr.com/research/oracle-switches-autopilot-turns-heat-mysql-database-service>

---

<sup>2</sup> For more on the history of MySQL, see: MySQL, Wikipedia. <https://en.wikipedia.org/wiki/MySQL>

---

<sup>3</sup> For more, see Holger Mueller, "Next-Gen Database Market Overview," Constellation Research, November 13, 2017. <https://www.constellationr.com/research/next-gen-database-market-overview>

---

<sup>4</sup> Holger Mueller, "The Universal Database Versus a Suite of Specialized Databases," Constellation Research, April 5, 2021. <https://www.constellationr.com/research/universal-database-versus-suite-specialized-databases>

---

<sup>5</sup> Holger Mueller, "There Is a HeatWave for Databases: Oracle MySQL Database Service Merges OLTP and OLAP," Constellation Research, June 3, 2021. <https://www.constellationr.com/research/there-heatwave-databases-oracle-mysql-database-service-merges-oltp-and-olap>

---

<sup>6</sup> Holger Mueller, "Oracle MySQL HeatWave Brings Machine Learning Inside the Database," Constellation Research, March 29, 2022. <https://www.constellationr.com/research/oracle-mysql-heatwave-brings-machine-learning-inside-database>

---

<sup>7</sup> The author uses the term Enterprise Acceleration to describe 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 a vice president and principal analyst at Constellation Research. He provides 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, chief technology officers, chief product officers, 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](https://constellationr.com/users/holger-mueller)  [linkedin.com/in/holgermueller](https://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](http://www.ConstellationR.com)

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

 [info@ConstellationR.com](mailto:info@ConstellationR.com)

 [sales@ConstellationR.com](mailto: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 | Denver | Ft. Lauderdale | New York Metro  
Northern Virginia | Portland | Pune | San Diego | Sydney | Washington, D.C.