

OFFERING OVERVIEW

Oracle MySQL HeatWave Brings Machine Learning Inside the Database

HeatWave Does Machine Learning the Right Way,
Gains Real-Time Elasticity, and Lowers TCO



Holger Mueller
Vice President and Principal Analyst

TABLE OF CONTENTS

Executive Summary..... 3

About Oracle MySQL HeatWave 4

Functional Capabilities 9

Analysis and Observations15

Recommendations.....17

Related Research19

Endnotes22

Analyst Bio23

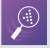


About Constellation Research.....24



EXECUTIVE SUMMARY

This Offering Overview examines the new capabilities of Oracle MySQL HeatWave that were added in late March 2022. 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.¹

Business Themes

-  Data to Decisions
-  Technology Optimization
-  New C-Suite

ABOUT ORACLE MYSQL HEATWAVE

Overview

MySQL became part of Oracle via the Sun Microsystems acquisition in 2009.² 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 in December 2020 with the addition of HeatWave, a cloud-native in-memory query accelerator designed to accelerate OLAP and other complex queries within the MySQL database.

Oracle continued its investments and innovations for MySQL, with the addition of machine learning-based automation to HeatWave with a component called MySQL Autopilot in August 2021. Autopilot brings management and performance automation capabilities to an open source cloud-native database without any additional charge. And in spring 2022, Oracle expanded the MySQL HeatWave service further, with the addition of HeatWave ML to the already available OLTP and OLAP capabilities inside HeatWave, automating the machine learning lifecycle and making it less of a data scientist domain. Oracle announced more innovation with real-time elasticity and doubled the amount of data processed in a HeatWave node to reduce costs, added support for pause and resume, and provided new benchmark results. The most prominent features of MySQL HeatWave ML are highlighted in this Constellation Offering Overview.

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—except for those offering the two cloud-native options—make their

Figure 1. Seven Market Trends Defining Next-Gen Databases



Source: Constellation Research

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 (expensive) memory. 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 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 can 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.
- 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 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. **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 AI/ML Conundrum

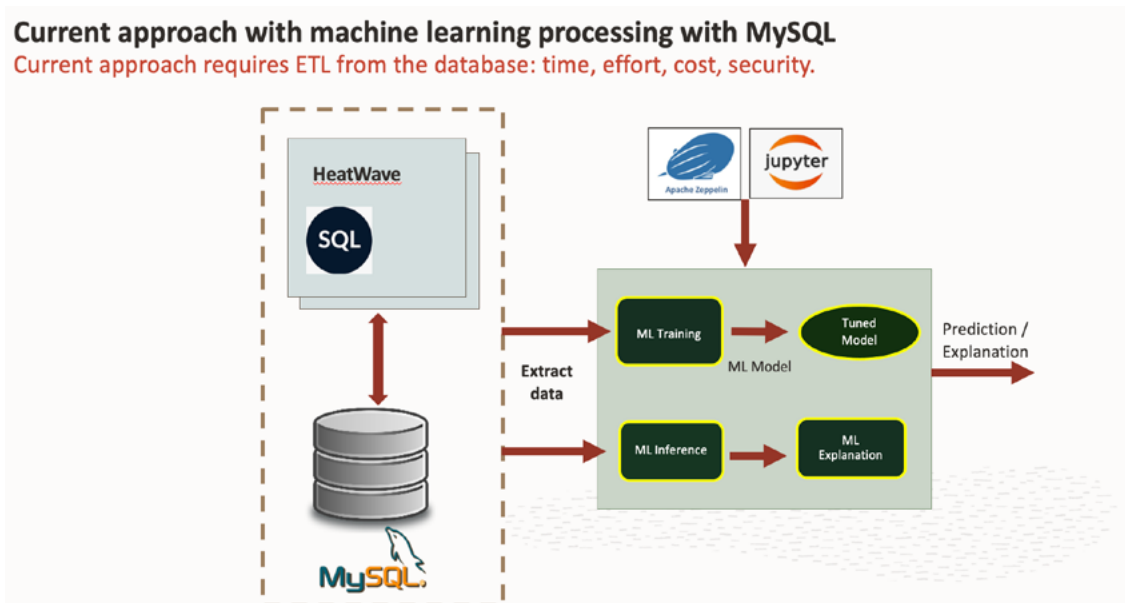
Traditionally, IT systems have been specialized, for a variety of reasons. Hardware capacity limitations have been a key constraint, leading to a segmented systems landscape that has deeply determined how IT is run. With the advances in hardware and the availability of Infinite Computing³ platforms, it is possible to blur these boundaries and overcome them for better automation of enterprises, which literally powers Enterprise Acceleration.⁴

MySQL Heatwave has already demonstrated that possibility by overcoming the boundaries between OLTP and OLAP and workload-aware automation through Autopilot. Now MySQL HeatWave is bringing AI/ML to the data in its database. This method has been tested and proven in the past with the “bring compute to the data” approach, which has been used in high-performance computing (HPC), operating applications at the edge, and more.

Traditionally, enterprises would have used an extract, transform, and load (ETL) offering to move the data out of their databases into a staging area from which they would build their AI/ML applications (see Figure 2). This approach has key drawbacks, the most prominent being that the data is always delayed, making real-time decisions powered by AI/ML impossible. Operating ETLs is a cumbersome, slow, often manual, and error-prone process that can leave the AI/ML automation “stranded” because it is not getting new data from the source systems—with possibly disastrous consequences for an enterprise relying on AI automation to guide consumer or sales behaviors. Moreover, moving data and the ML model out of the database leads to data proliferation and security vulnerabilities. The result is a time-consuming, expensive, and less than optimally secure path to ML insights and ML-powered applications.

But what if a design could bring the ML to the data, eliminating the need for ETL processes and enabling faster and close to real-time ML usage? That’s what the MySQL HeatWave team has achieved.

Figure 2. Traditional ML Operations—ML Outside the Database



Source: Oracle

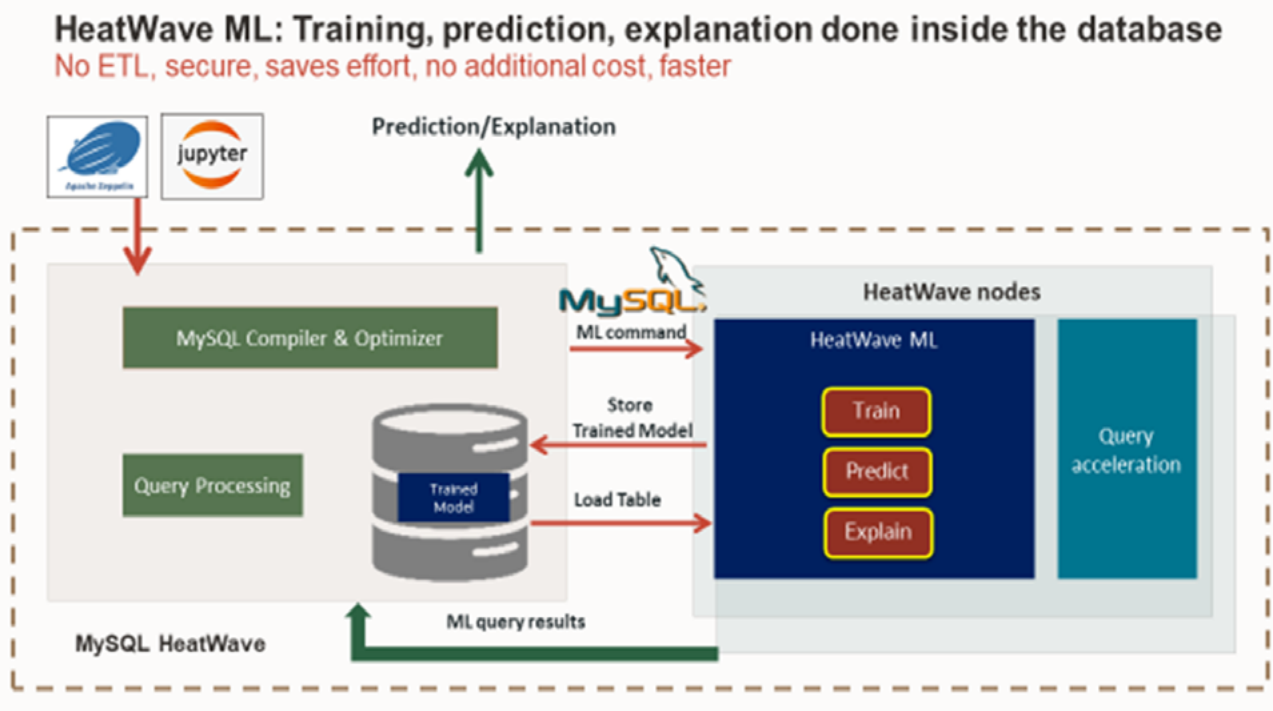
FUNCTIONAL CAPABILITIES

HeatWave Brings AI to the Data

With its latest release, MySQL HeatWave is ending the separation of database data and data used to model ML applications. ML models can be built on the HeatWave nodes with HeatWave ML, which automates the training, inference, and explanation of the ML models. The trained models are stored in the database and accessed via the MySQL compiler and optimizer. Users can access the models via Apache Zeppelin and/or popular Jupyter notebooks (see Figure 3).

The result is a tightly coupled, closely integrated data and ML combination that enables faster and cheaper automation with ML. Moreover, it moves ML execution and modeling closer to real time, because the elimination of ETL processes enables ML to directly execute on the data stored in the OLTP area of HeatWave and enhances security, since the data does not leave the database. There are no data movers or connectors involved—as in this case, in-database literally means inside the MySQL HeatWave database service.

Figure 3. HeatWave and Its Embedded ML Capabilities



Source: Oracle

Fully Automated ML Training With HeatWave


Creating ML models requires both the domain experience and expertise of a data scientist. Typically, a data scientist performs the task—the challenge for enterprise is that there are not enough data scientists available for the number of business problems that ML can address, solve, and mitigate.

Oracle addresses the shortage of skilled ML professionals by fully automating the training for ML models inside HeatWave. Practically, HeatWave uses ML to create ML models, employing automation via MySQL Autopilot.

In detail, this is how HeatWave creates ML models (see Figure 4):

- 1. Automated preprocessing.** HeatWave ML automates the necessary preprocessing steps such as data cleansing and normalization. This equates to huge cost and performance savings compared with any manual approach to data preparation.
- 2. Algorithm selection.** Next, HeatWave ML automatically identifies the top algorithms suitable to the ML problem at hand. These meta learned proxy models allow only one pass at every pipeline stage and determine the next processing steps as well as feature selection and sampling.
- 3. Feature selection.** HeatWave ML automatically identifies the relevant features for running a successful model, by determining which columns in the database are the more important ones for model success.

Figure 4. Traditional ML Operations—ML Outside the Database



Source: Oracle

4. **Adaptive sampling.** Next, HeatWave ML selects the appropriate sample size and extent to use for training the model. Remarkably, HeatWave identifies potential imbalances in the sampling and adjusts for them.
5. **Hyperparameter optimization.** Model success stands and falls with optimization of the model's hyperparameters. HeatWave ML has a novel gradient-based search space reduction technique that provides a high degree of parallelism without compromising accuracy while determining the hyperparameter.
6. **Model explainer.** HeatWave ML has integrated explanations with model training. As such, all models generated by HeatWave ML can be explained (see Figure 5).
7. **Prediction explainer.** HeatWave ML is also able to provide prediction explanations. Since this is done at the time of training, the system can glean the necessary features and metadata from the training data and store it in the trained model. For explaining the predictions, HeatWave ML does not require training data. This is a valuable differentiator.
8. **Tuned model.** And finally HeatWave ML creates a tuned ML model, ready for use. This tuned model is stored inside the MySQL database and is secured by the database security mechanisms.

Finally, in the spirit of Oracle's ambition to make all software autonomous, HeatWave ML provides automated tuning and training of models, able to accelerate each pipeline stage with unique parallel algorithm considerations. HeatWave ML scales with the number of available clusters, making it faster, cheaper, and more secure as a single database service than other, more traditional offerings.

Not only is the training with HeatWave ML fully automated but it is also very fast. The advantage of having a very efficient training process is that models can be retrained and kept up to date with changes in the data. This improves the accuracy of the model for making more accurate predictions.

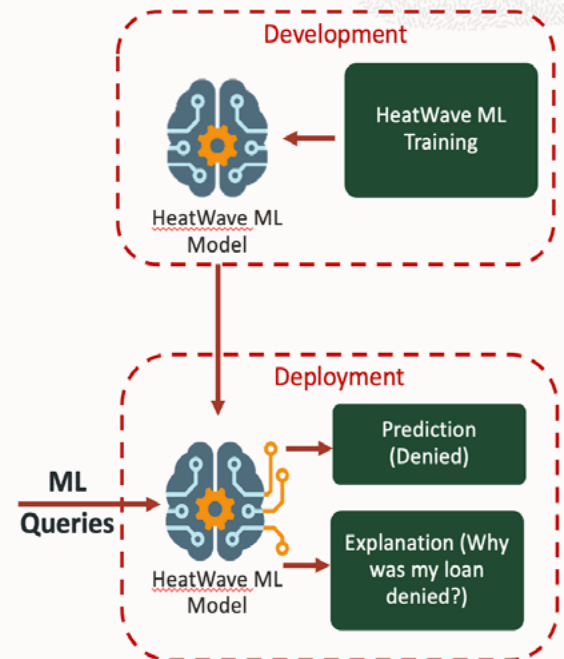
Figure 5. All Models Created in Heatwave ML Can Be Explained

ML Explanation is important for enterprise workloads

- Explanations are integrated with the training pipeline
- Model agnostic techniques can explain any HeatWave ML model

Helps:

- **Regulatory compliance:** may imply 'right to an explanation' for algorithms affecting users
- **Fairness:** by allowing validation that predictions are unbiased
- **Repeatability:** ensures that small changes in input do not lead to large changes in the explanation
- **Causality:** allows verification that only causal correlation between features and predictions are selected
- **Trust:** interpretable explanations encourage machine learning based prediction



Source: Oracle

Real-Time Elasticity Comes to HeatWave

Elasticity of computing resources is critical for cost-effective operation of software workloads in the cloud. Modern software needs to be able to scale both up and down, thus consuming different amounts of computing resources and, in the process, scaling costs to match consumption.

The dirty secret behind elasticity is that it often is achieved via manual processes and is one way: up. But manual processes can be slow and are expensive and often error-prone, leaving the database inaccessible for some time. MySQL HeatWave addresses the elasticity challenge, by parallelizing the resize process for more capacity, making downtime an experience of the past (see Figure 6).

MySQL HeatWave can provision additional nodes dynamically, without requiring any downtime and any manual steps on the part of the user. Importantly, the real-time elasticity capabilities of MySQL HeatWave do not limit key database operations (queries, data manipulation, loading, and so forth) during resizing. They all can happen on the cluster during the process. The process is also flexible, enabling

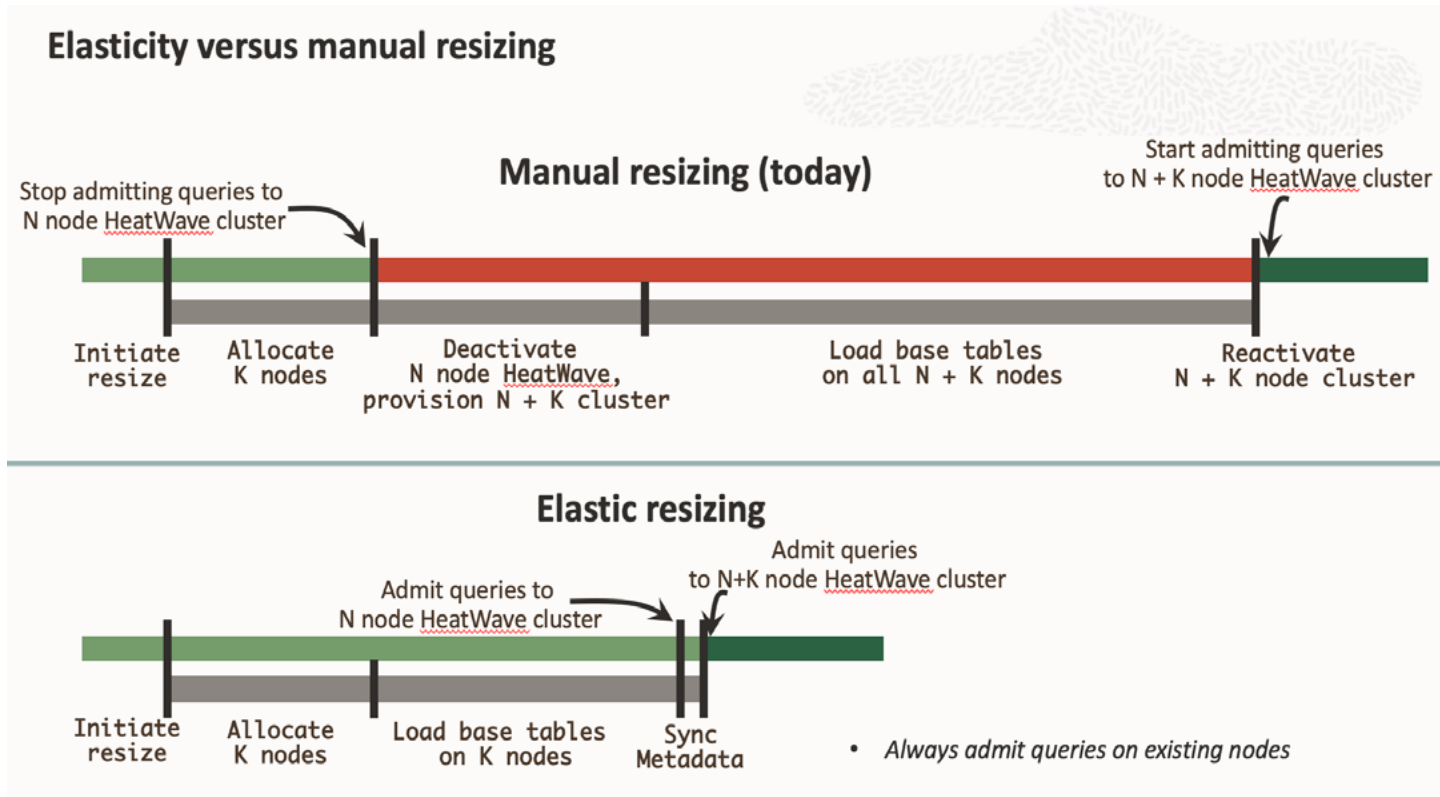
users to upsize or downsize to any number of nodes as needed. And in the tradition of elegant architecture approaches' being applied with MySQL HeatWave, the resizing requires minimal data movement, because data is loaded at an object level.

The result is excellent elasticity performance, with resize times dropping exponentially when nodes are added and pauses caused by table loading being measured in the microsecond range. Furthermore, at the end of the elasticity operation, the data across the various nodes of the cluster is automatically balanced.

HeatWave Becomes (Even) More Cost-Effective

Oracle has a very long tradition of reducing IT costs—in fact, the company was founded with total cost of ownership (TCO) in mind. So it comes as no surprise that HeatWave follows the path of other Oracle products by reducing its footprint and the related operational costs.

Figure 6. The Benefits of Elasticity Over Manual Resizing



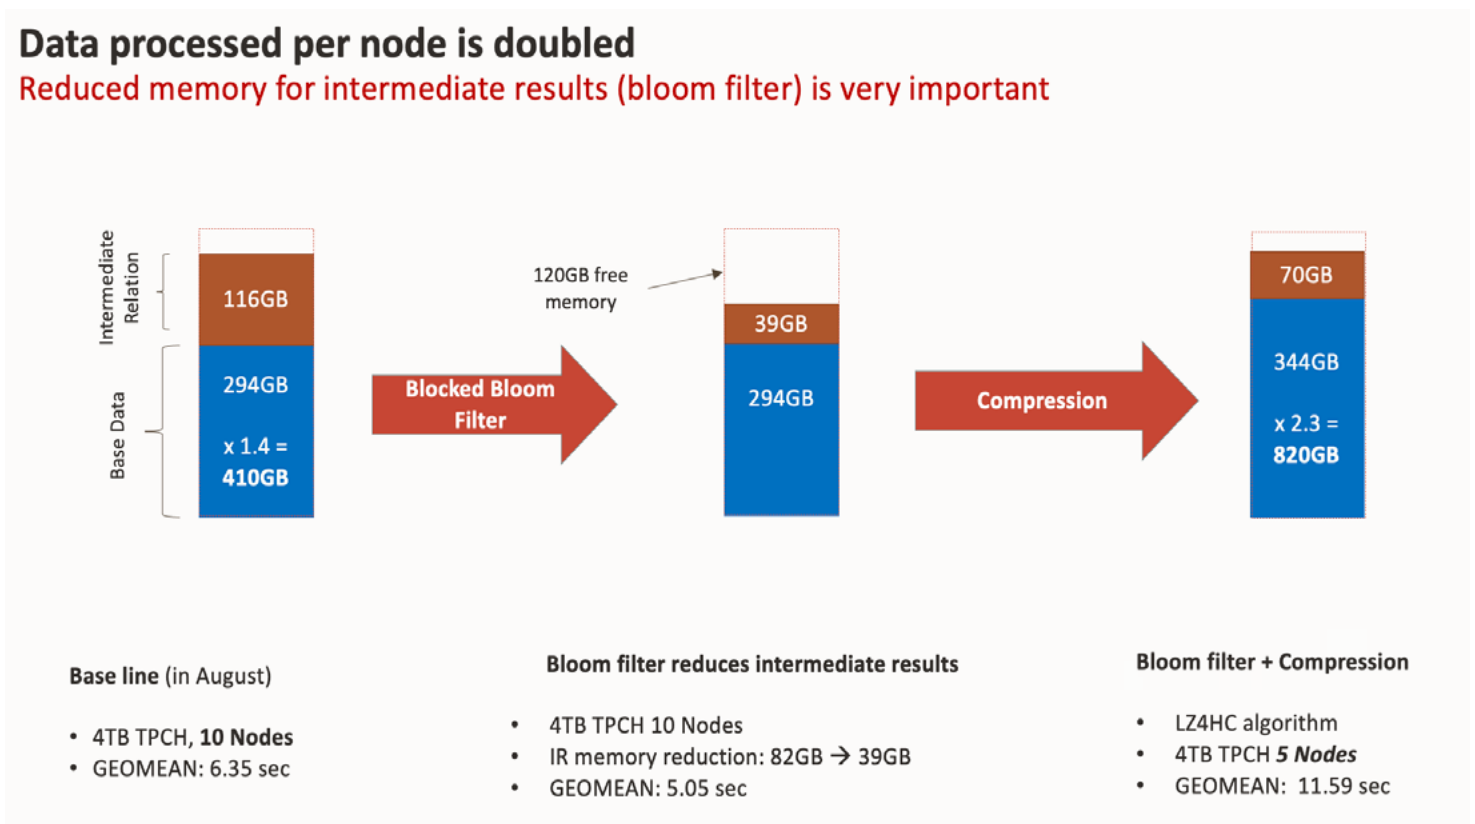
Source: Oracle

Among a plethora of improvements, the one yielding the most in terms of reducing the HeatWave footprint has been the replacement of standard bloom filters with blocked bloom filters (aided by parallelized advanced vector instructions [AVXs]), which turned out to be three times as efficient as the standard bloom filters. Lookup times, for example, were reduced from 45 nanoseconds to 15 nanoseconds.

What is remarkable is the speed at which the HeatWave team has been able to reduce the database’s footprint, practically doubling the workloads that can run on the same cluster size compared with the version released in August 2021 (see Figure 7).

Oracle also introduced a new “pause-and-resume” function, which enables customers to pause HeatWave to save additional costs. When users invoke the “resume” function, both the data and the statistics needed for MySQL Autopilot are automatically reloaded into HeatWave.

Figure 7. The Impressive Cost Reduction of MySQL HeatWave



Source: Oracle

ANALYSIS AND OBSERVATIONS

Constellation sees the following strengths and weaknesses for the latest Oracle MySQL HeatWave release (see Table 1):

Strengths

- **It offers a unique combination of OLTP, OLAP, and ML.** Enterprises need to accelerate and become more agile, and that starts with decision management. Insight to action has long eluded CxOs, because MySQL cloud databases often require ETL to a separate analytical database, thereby delaying insights that are, at this point, based on stale data. MySQL HeatWave facilitates insight to action in real time, enabled by its backward compatibility, high performance, and low cost, and with subsecond transactional changes reflected in analytical insights.
- **It provides a high degree of automation with MySQL Autopilot.** The MySQL Autopilot capabilities make HeatWave a highly automated database. In times when there is a dearth of trained resources, this is a key differentiator for CxOs concerned about the uptime of their next-generation applications.
- **ML operates directly on the data.** As its latest major innovation, HeatWave is operating ML inside the database, effectively bringing the ML to the data. This lowers TCO, reduces the need for manual work and ETL, and accelerates the ML model provision and execution that enterprises need.
- **It is another testament to Oracle's TCO-focused organizational DNA.** By adding ML in the database with MySQL HeatWave with no additional license cost, Oracle delivers an innovation that is another proof point of its corporate DNA: It is all about lowering the TCO of running IT—in this case with MySQL HeatWave. MySQL HeatWave follows in those footsteps, providing leading category performance while offering lower costs via better technology and engineering, delivering several innovative capabilities added in the last release. CxOs can count on Oracle to provide better price/performance, often within short timeframes—in this case, three new major releases for HeatWave in less than 15 months.

Weaknesses

- **It is 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 MySQL HeatWave, which new customers should test thoroughly for both scale and functionality. However, HeatWave is rapidly maturing as a service, and it is one of the fastest-growing services on Oracle Cloud Infrastructure (OCI).
- **It is available only in OCI.** This may well be a moot point in the short term (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. Oracle can address this by making HeatWave available on other cloud platforms. In fact, Oracle Executive Chairman and Chief Technology Officer Larry Ellison mentioned in the Q3FY2022 earnings call that MySQL HeatWave will be available on AWS. Enterprises want to avoid cloud lock-in and therefore favor offerings that support/are available in multiple clouds.
- **It is built. How many will come?** 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 machine learning workloads will be run best and first by MySQL HeatWave and which independent software vendors (ISVs) will take up the new service. Note that no changes are required to run MySQL applications on HeatWave.

Table 1. MySQL HeatWave Strengths and Weaknesses

STRENGTHS	WEAKNESSES
<ul style="list-style-type: none">• Is a unique combination of OLTP and OLAP + ML.• Is highly automated with MySQL Autopilot.• Operates ML directly on the data.• Is another testament to Oracle's TCO-focused organizational DNA.	<ul style="list-style-type: none">• Is a new offering.• Is available only in Oracle Cloud Infrastructure (for now).• It is built. How many will come?

Source: Constellation Research

RECOMMENDATIONS

Constellation offers the following recommendations regarding MySQL HeatWave, MySQL Autopilot, and HeatWave ML:

- 1. Enable Enterprise Acceleration.** 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. Not having to change code to adopt MySQL HeatWave is a major benefit and will help the adopting enterprise move faster and accelerate with automation provided by MySQL Autopilot.
- 2. Support and then practice insight to action and accelerate ML adoption.** Enterprise decision-making has been hampered by the delay of OLTP data availability in OLAP systems. MySQL 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 enables real-time insight-to-action best practices. And by bringing ML to the data with HeatWave ML in a cost-efficient, automated way, HeatWave accelerates ML adoption.
- 3. Evaluate MySQL HeatWave now.** Customers using MySQL or MySQL-compatible databases should assess the immediate benefits of adopting MySQL HeatWave. Especially with the automation now available with MySQL Autopilot and in-database machine learning, those benefits are so compelling that existing Oracle MySQL customers, regardless of on-premises or cloud deployments, should immediately adopt the new offering. The new ML features, and the reduced footprint, make HeatWave even more attractive. 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. Consider the impact of native ML capabilities.** MySQL HeatWave has made an already-attractive offering even more attractive. As if the combination of OLTP and OLAP workloads in a single MySQL database and the recent addition of MySQL Autopilot were not attractive enough, Oracle has managed to make the offering even more compelling with HeatWave ML and its in-database machine learning capabilities. Enterprises need to tap into the powers of ML running close to their transaction data to build the ML they need in order to excel in the current competitive markets.

5. **Conduct a cost/benefit analysis for a potential switch to Oracle MySQL HeatWave.** Non-MySQL customers need to run the numbers. Oracle has achieved quite an engineering feat with MySQL 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 to copy the Oracle technological approach to bringing OLTP, OLAP, and ML together inside one database, it requires serious talent and support to provide a long-lasting, future-proof, and trusted platform.

6. **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 MySQL 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 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 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 9, 2022. <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>

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 9, 2020s. <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, February 16, 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 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 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 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>

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 GraalVM 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 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>

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

³ For more on Infinite Computing, see: Holger Mueller, "The Era of Infinite Computing Triggers Next-Generation Applications," June 1 2018. <https://www.constellationr.com/research/era-infinite-computing-triggers-next-generation-applications>

⁴ Mueller uses the term Enterprise Acceleration to refer to 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](https://www.constellationr.com/users/holger-mueller) [in linkedin.com/in/holgermueller](https://www.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.