

AUTOPILOT DRIVES MYSQL DATABASE SERVICE DIFFERENTIATION

EXECUTIVE SUMMARY

Time-to-value and time-to-action are two key metrics of success in the digitized economy, both measuring access to actionable intelligence. This intelligence is only useful when all data sources are appropriately transformed and analyzed.

MySQL, an open-source relational database, gained widespread popularity in companies of all sizes in the early 2000s. Today, MySQL instances are the back-end of many applications running across enterprises of all types, from content streaming providers to social media platforms to the world's leading financial services companies. Countless invaluable data exists in these MySQL instances – data that, until recently, was difficult to aggregate holistically across a business.

Oracle introduced MySQL Database Service with HeatWave in late 2020 to simplify the transformation of MySQL data into intelligence. HeatWave is a massively parallel, high-performance, in-memory query accelerator for Oracle MySQL Database Service that enables real-time analytics on the data residing in MySQL databases. HeatWave is designed to scale to hundreds of nodes running in the cloud and optimally works with Oracle Cloud Infrastructure (OCI).

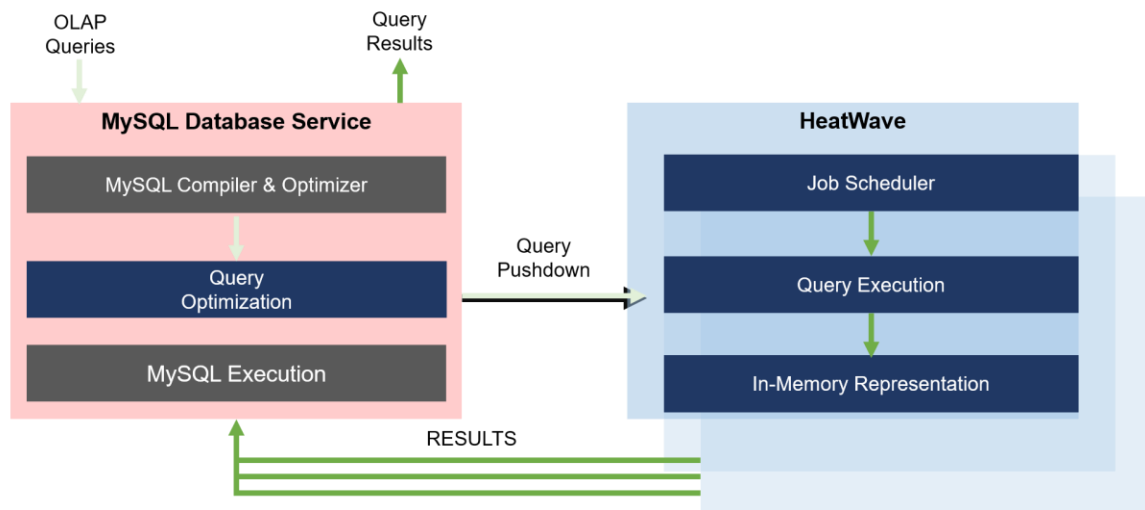
The innovation of MySQL HeatWave, however, didn't end with its initial release. Oracle recently upgraded MySQL HeatWave with Autopilot, focusing on machine learning (ML)-based automation to speed up its customers' time to obtain valuable business insights and improve database optimization and performance. This research brief will drill down on the latest innovations in MySQL HeatWave and explore how customers of all sizes can benefit from utilizing Autopilot.

HEATWAVE – AN INTRODUCTION

As a refresher, Oracle MySQL HeatWave is a cloud-native service available in OCI to enable the real-time analytics of traditional MySQL environments supporting online transaction processing (OLTP). The key to this description of HeatWave is “real-time.” Data generated in a customer's MySQL environment automatically synchronizes in HeatWave. As a result, analytics or other queries running against HeatWave analyze

the most current data without the need for an extract, load, and transform (ELT) process or other tools.

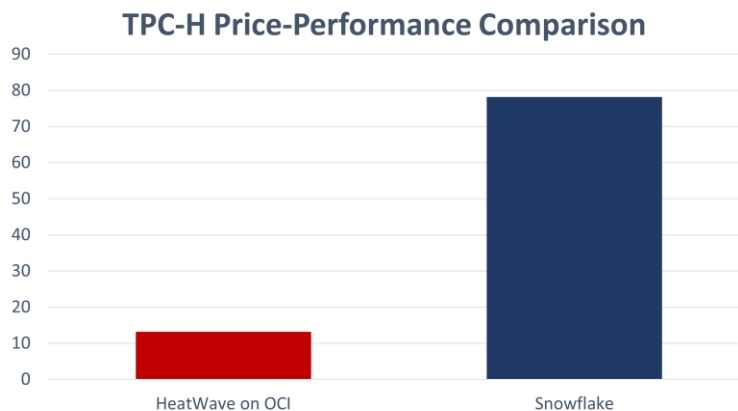
FIGURE 1: HEATWAVE REAL-TIME ANALYTICS



Source: Moor Insights & Strategy

While many cloud-based offerings such as Snowflake offer analytics services, HeatWave’s unique architecture can lead to significant cost and performance advantages. Because of these advantages, MySQL HeatWave has been well-received, and customers from other clouds are migrating to OCI.

FIGURE 2: HEATWAVE PRICE – PERFORMANCE LEADERSHIP



Source: Moor Insights & Strategy

As you can see in the above comparison graph, the HeatWave service has a significant price/performance advantage over Snowflake in TPC-H benchmark testing (which Oracle has made [available](#) for the public to use in testing customer-specific datasets). For TPC-H, HeatWave demonstrated 7x better performance at one-fifth the price compared to Snowflake's "pay-as-you-go" pricing. That's an amazing 35x price/performance advantage for HeatWave.

It is not hyperbolic to say that Oracle's innovation was on full display in HeatWave as these price and performance advantages showed similar results against other cloud-based analytics providers.

THE HEATWAVE INNOVATION CYCLE

Oracle's release of HeatWave marked a wave of innovation in MySQL support. The ability for organizations to run complex OLTP and real-time analyses of MySQL environments so much faster and cheaper than the competition was a significant feat, and poses a challenge to existing cloud database vendors.

One could argue that this feat is significant because HeatWave represents a true, cloud-architected solution from a company that made its fortunes in traditional software development and support.

As one would expect from a cloud-native service, Oracle continues to drive functionality in HeatWave, with a new component targeting the automation of initial deployment, data placement, query performance, and failure handling of MySQL HeatWave. As with any cloud computing service, this deployment model enables customers to automatically realize performance and management benefits without disrupting existing MySQL environments or making any changes to applications.

Moor Insights & Strategy sees this approach from Oracle as quite compelling. It combines its vast knowledge of the data management space with its experience in automating enterprise cloud services.

INNOVATION GROUNDED IN THE REAL WORLD

A company accustomed to developing, delivering, and improving a product over its lifecycle understands the balance of delivering what the market wants in the short-term and what it needs in the longer-term. These two requirements are sometimes at odds with one another. But understanding this tension and building out a product

development and release strategy enable a company with 40 years of experience to succeed in delivering innovative cloud services. And that is what Oracle has done with HeatWave.

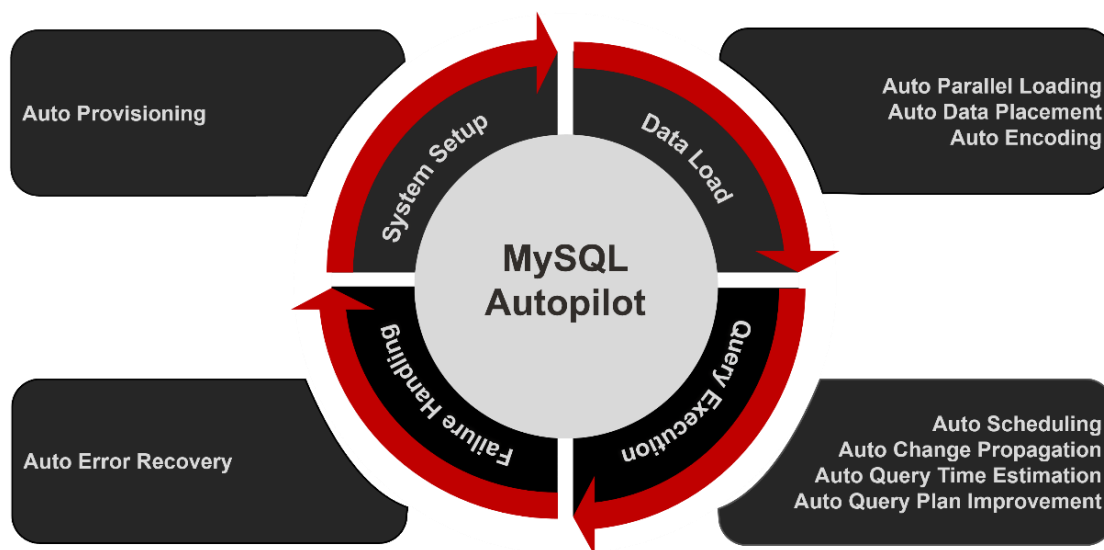
The next generation of the HeatWave cloud service demonstrates this approach by striking this balance. In this latest release, Oracle has built-in several ML-driven tools that simplify the deployment, use, performance, and management of HeatWave while building a management framework that can deliver a state of automation unavailable from other competitive cloud services. MySQL Autopilot embodies that balance.

AUTOPILOT – AUTOMATING DATA LIFECYCLE MANAGEMENT

Database administrators (DBAs) and other professionals tasked with managing data environments can attest to the challenge of maintaining data management platforms – from initial creation to the constant tuning required for optimal performance to staying up to date on patching. Further, this is made even more challenging by the proliferation of databases like MySQL across the enterprise.

Oracle appears to see the answer to these challenges through the activation of automated operations. At the heart of Oracle's response is MySQL Autopilot, an automation engine driven by deep analytics and ML to power nine functions across four areas of data management.


FIGURE 3: MYSQL AUTOPILOT



Source: Moor Insights & Strategy

Autopilot extrapolates statistics from MySQL environments on a per-instance basis. These instance-specific statistics feed ML models that, in turn, tune MySQL HeatWave. These capabilities remove much of the manual intervention required by DBAs and IT professionals to not only maintain but optimize a MySQL environment.

FIGURE 4: DEEP ANALYTICS AND ML DRIVE AUTOMATION IN AUTOPILOT

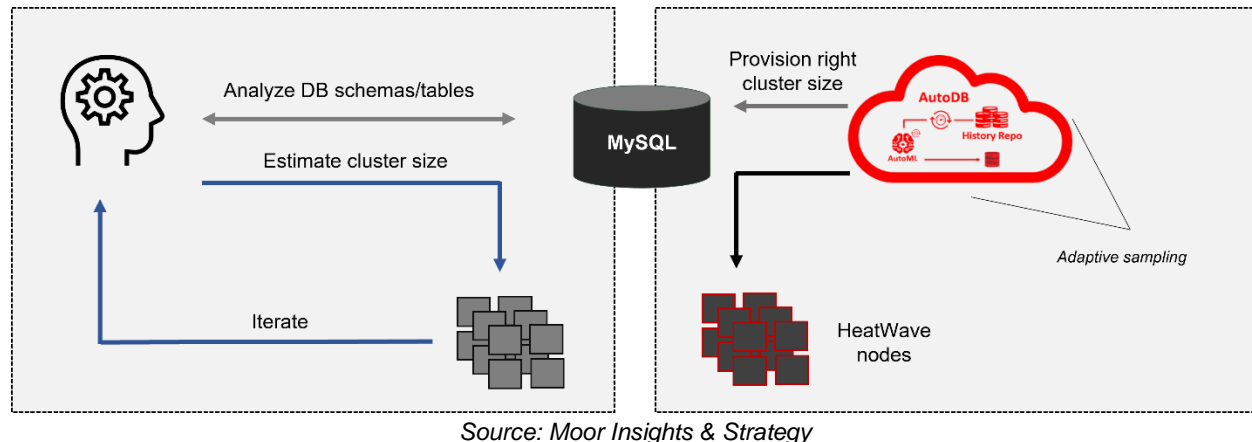


Source: Moor Insights & Strategy

A good example of the efficiencies gained through MySQL Autopilot is its Auto Provisioning capability. Estimating the size and shape of a database cluster can be frustrating for a database professional. It is equal parts science and guesswork. DBAs look at database schemas and tables, consider use, performance, and data growth and estimate the required resources for optimal performance. And from this point forward, right-sizing that cluster is a never-ending iterative process, which takes many hours and often leads to the overprovisioning of hardware.

MySQL Autopilot resolves this system setup task. With Auto Provisioning, an ML algorithm looks at a small slice of table data and predicts the amount of memory required to determine the right size of the HeatWave cluster. Even with this small sampling (under 0.1%), Autopilot can predict the right size of the cluster with a high degree of accuracy (>97%).

FIGURE 5: AUTOPROVISIONING PROCESS



Source: Moor Insights & Strategy

The impact of three other capabilities offers examples of benefits of Autopilot:

- **Auto Data Placement** optimizes the placement of columns in memory and key placement based on customer- and data-specific query history. The system recommends the optimal columns and also predicts the expected improvement from this change
- **Auto Query Improvement** learns from the queries executed in the past and improves the SQL query plan of the new incoming queries.
- **Auto Scheduling** manages the query pipeline across online analytical processing (OLAP) and OLTP workloads. Through this capability, shorter queries, typically OLTP, are scheduled to execute ahead of long analytical queries (reports, for example), reducing the overall wait time for queries. Most other cloud databases use the First In, First Out (FIFO) approach for scheduling.

While each of these three capabilities is compelling on its own, it's the symbiotic way that these services work together that drives an ever more significant benefit to data managers and IT organizations. Without these capabilities, DBAs face an impossible task of maintaining an organization's data environments – constantly tuning but never fully optimizing.

Finally, it's important to note that these capabilities are customer- and database-specific. This is significant because it ensures the data feeding ML models and algorithms is tailored for that specific environment. Cluster-wide or cloud-wide models that can apply refinements based on the aggregate performance of other instances could likely end up negatively impacting the performance and cost of a customer's database.


MYSQL HEATWAVE SCALE-OUT DATA MANAGEMENT

The other improvements made to MySQL HeatWave are the architectural support and the loading and reloading of data. In the initial release of HeatWave, data loads could be somewhat time-consuming – as any DBA would expect when loading and reloading data into tables.

Because of the highly parallelized architecture consisting of scale-out nodes based on AMD EPYC CPUs and object store in OCI, data is loaded and stored encrypted in an in-memory format in the object store. When data needs reloading into the HeatWave cluster, it is read from the object store at near object store bandwidth. Not only does data load at high speed in a highly parallelized method, but it is also highly partitioned, allowing for granular levels of data management.

This faster reload into HeatWave improves the availability of HeatWave cluster up to 100x for 10TB data size

FIGURE 6 – SCALE-OUT DATA MANAGEMENT IN HEATWAVE



Source: Moor Insights & Strategy

SCALE-OUT ACROSS MORE NODES

In addition to the capabilities detailed in the previous sections, Oracle implemented several other enhancements of note. One improvement is the support for larger data sizes (up to 32TB), more than doubling the number of nodes (from 24 to 64) and improving scalability by 20%. This capability will enable HeatWave to support the demands of medium to large companies for both reporting and analytics.

Another area of improvement is HeatWave's ability to drive faster times for complex OLTP queries. While the initial release of HeatWave targeted analytics, the less complex OLTP query times executed at par in performance with other cloud services, though cost savings were recognized. MySQL HeatWave now accelerates many other constructs, which often appear in mixed workloads. MySQL Autopilot's Auto Scheduling service prioritizes simple OLTP queries, driving faster response times in mixed workload environments relative to the competition.

Security for data-in-flight is another area of improvement that customers can expect with Oracle's latest release of HeatWave. This capability, combined with OCI's inherent security and AMD EPYC's security features, should give customers extreme levels of confidence.

Lastly, HeatWave accelerates all the 94 TP-CDS queries supported by MySQL in the latest version of the HeatWave service.

ADDING IT ALL UP – HOW MYSQL DATABASE SERVICE WITH HEATWAVE AND AUTOPILOT FURTHER BENEFITS THE DATA CONSUMER

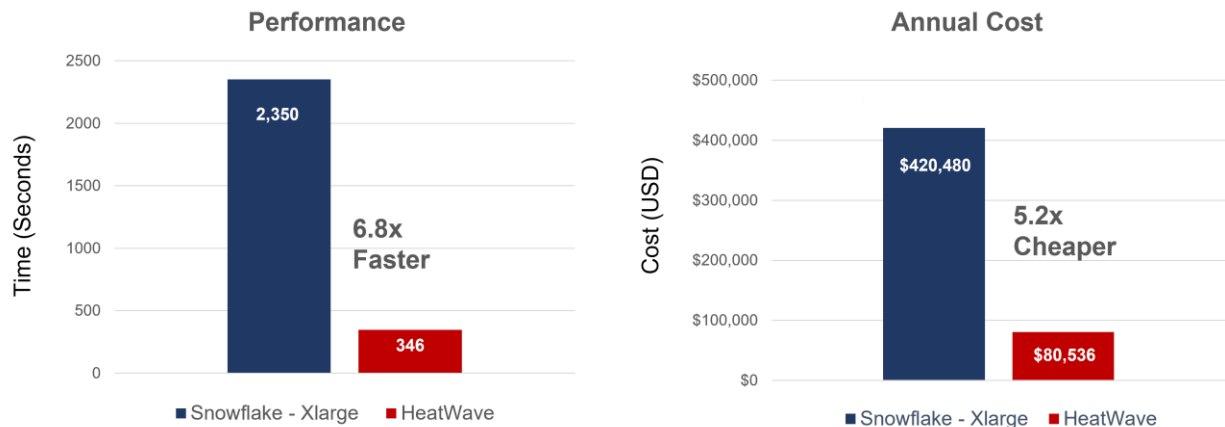
As any IT professional can attest, product and service upgrades, features, and capabilities are only relevant if the customer can derive measurable benefits relative to alternatives. And this is where Moor Insights & Strategy finds MySQL HeatWave to be compelling.

In the initial release of HeatWave, Oracle demonstrated strong performance and price/performance advantages relative to the competition. Oracle's latest release of HeatWave increases that gap by focusing on several areas, all of which should deliver additional business value to data consumers:

- **Achieving faster time-to-value** – This is a return on investment (ROI) component of every IT solution that is sometimes undervalued as a metric. Many vendors tend to obfuscate this through marketing literature and slogans. However, ROI is critical in determining value. Oracle has implemented new MySQL Autopilot features like auto data placement and auto query planning, resulting in this faster time-to-value and increased ease-of-use.
- **Improving performance** – As previously demonstrated, HeatWave showed a strong performance advantage relative to other cloud services. Oracle has achieved more robust performance and price/performance results relative to

Snowflake and other cloud services through a combination of its improvements to the HeatWave architecture, many of the automation capabilities built in to Autopilot and support for all TPC-DS queries.

FIGURE 7 – HEATWAVE V SNOWFLAKE



Source: Moor Insights & Strategy

In the above chart, the benefit of HeatWave's improvements is genuinely realized by a near 7x performance advantage and a 5x price advantage over Snowflake. While these numbers may look too good to be true, Oracle took the unprecedented step of publishing all its testing parameters on [GitHub](#) for customers to download and test in their own environments. With this repository, customers can compare HeatWave against other hyperscale cloud service providers anytime, anywhere.

- Automating operations** – Another area where Oracle made substantial progress is the automation of operations. Perhaps put more cynically, the automation of the mundane. Moor Insights & Strategy spoke with data management and IT professionals, and operations is a consistent leader on the list of jobs requiring the most time but delivering the least satisfaction and value to the organization.

There is a cost and time component to IT and database operations that are quantifiable. However, there's also an employee satisfaction element that is equally important, though far less quantifiable. In a tight labor market, organizations should create challenging and stimulating environments for their employees. Offloading many of these operational tasks to MySQL Autopilot can certainly help achieve that goal.

- **Increasing resiliency** – Great products delivering great value means nothing if they are neither reliable nor secure. With the next generation of HeatWave, Oracle appears to have achieved both through Auto Error Recovery and the addition, as mentioned earlier, of securing data-in-flight. These capabilities complement the resiliency, security, and reliability inherent in Oracle Cloud Infrastructure, such as automatic testing and applying of updates and patches. Though technically falling under the "operations" category, these capabilities directly impact environment resiliency.

It is fair to say that Oracle has positioned itself firmly as a leader in the MySQL space with HeatWave plus Autopilot based on Moor Insights & Strategy's evaluation of claims, testimonials, and interviews with data management professionals.

IN CLOSING

Data management and the ability to extract the most from data dispersed across enterprises is key to success in any digitization effort. Moor Insights & Strategy believes that organizations quickly transforming and acting upon the glut of data generated every minute of every day are the organizations that will lead in the future.

The fundamental challenge in digitization is transforming and analyzing vast amounts of data stored in (up to) thousands of MySQL instances across an enterprise. While many database cloud services offer analytics services, MySQL HeatWave is the first to enable real-time analytics of MySQL environments without the need for costly ETL functions to a second analytical database. It's the first cloud service to combine OLTP + OLAP + ML automation for MySQL applications at rates of performance that eclipse everything from Snowflake to Alibaba and anything in between.

With MySQL HeatWave, Oracle demonstrated significant and verifiable price/performance advantages over the biggest cloud services, including Snowflake. An initial review of HeatWave, including its design, functionality, and performance, can be found [here](#).

In Oracle's latest release of the MySQL HeatWave service, the company seems to have focused on the onramp to this service, simplifying the use and management of HeatWave through Autopilot.

There are many reasons an organization chooses a database cloud service – cost, reliability, time in the market, relationships and cloud gravity (the number of services used from a cloud service provider). However, Moor Insights & Strategy sees the cloud services market evolving, whereby IT organizations are starting to consume services from multiple cloud providers based on capabilities, cost and other factors.

Moor Insights & Strategy sees Oracle Cloud Infrastructure as a strong player in the cloud services space and Oracle playing a disruptive role in the MySQL arena. As a result, we consider Oracle MySQL Database Service with HeatWave and Autopilot as a "must have" for organizations looking to derive the most from their MySQL environments.

For more information on MySQL HeatWave, visit www.oracle.com/heatwave

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

[Matt Kimball](#), Vice President and Principal Analyst at [Moor Insights & Strategy](#)

PUBLISHER

[Patrick Moorhead](#), Founder, President, & Principal Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

This paper was commissioned by Oracle. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2021 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.