

MySQL *HeatWave* DELIVERS AUTOMATION WITH AUTOPILOT

.....
An Open-Source Cloud Data Warehouse Service

ORACLE OFFERS A REMARKABLE CLOUD DATABASE SERVICE named *MySQL HeatWave*, aiming to provide distinctive performance and cost advantages for open-source database projects.

Released this summer, *MySQL HeatWave* is the first MySQL implementation with a built-in, highly parallel query accelerator, making it a viable option for data warehouses on Oracle Cloud Infrastructure (OCI), as well for the transaction processing workloads it has long handled.

MySQL HeatWave is much simpler to use and operate than the equivalent AWS product set. Customers using AWS native products need a more complex architecture and four separate products: *Amazon Redshift* for OLAP; *Amazon Aurora* for OLTP, and *Glue* and *Kinesis Firehouse* to move and transform (ETL) data between them.

MySQL HeatWave includes Oracle's new *Autopilot*, which employs machine-learning, to **automate the building and operation of databases**. With *Autopilot*, **optimization of query performance is further accelerated and simplified, eliminating another significant measure of effort and complexity for the database customer**.

MySQL HeatWave is aimed at a different set of requirements than Oracle Autonomous Database (ADB). *MySQL HeatWave* supports open-source applications such as *Wordpress*, *Drupal* and *Magento* — as well as targeting different buyers and developers by also supporting PHP, Python, Ruby, Go, Rust and NodeJS. By contrast, ADB supports Oracle Fusion SaaS Applications, the APEX low-code application development environment and JSON data.

While *MySQL HeatWave* is not an alternative to Oracle Autonomous Database (ADB) for high-end, enterprise requirements, it is not limited to the low end, either. In its current version, *MySQL HeatWave* supports configurations as large as 64 nodes — enough scale to handle significant workloads against tens of terabytes of data.

In 30 TB TPC-H benchmark results released by Oracle, *MySQL HeatWave* demonstrates a large performance and price/performance advantage against *Amazon Redshift*, *Amazon Aurora* and *Snowflake* on AWS (test results and code are available [here](#)). ●



WinterCorp
www.wintercorp.com

.....
SPONSORED
RESEARCH
PROGRAM

Methodology



PURPOSE AND METHODOLOGY FOR THIS REPORT

This *WinterCorp Research Note* covers the recent announcement of Oracle *MySQL HeatWave* and its implications for database customers. In developing this report, WinterCorp drew on its own independent research and experience, interviewed employees, attended Oracle events and analyzed Oracle documentation and literature. Oracle was provided an opportunity to comment on the paper with respect to facts, in its capacity as the sponsor of this research. WinterCorp has final editorial control over the content of this publication and is solely responsible for any opinions expressed.



SINCE THE 1990s, developers have turned to MySQL and other open-source software for an inexpensive approach to projects that require only the basic, standard database capabilities such as defining, storing, retrieving and updating data in relational tables. This approach has proven successful for a range of projects, most frequently those that are developer-led and that do not feature particularly demanding or complex database requirements.

MySQL was considered a good choice for many such applications, but not for projects that look much like a data warehouse. In a data warehouse, there are typically complex queries that must touch a lot of data. This can take a very long time, particularly as databases get larger and queries become more complex. As a result, MySQL customers typically would use a special purpose analytics database for data warehousing with all of the drawbacks and cost of having to ETL data across databases and manage two database environments — including increased security risks and surface area exposures, particularly while data is in flight.

With *MySQL HeatWave*, Oracle has changed that. *MySQL HeatWave* can parallelize across as many as 64 nodes (servers) to accelerate data warehouse queries with no change to the queries. That, alone, is a breakthrough that significantly expands the range of applications suited to MySQL. In addition, *MySQL HeatWave* is integrated with MySQL Database Service; customers don't have to ETL and manage a second database, avoiding management headaches and reducing overall costs by running OLTP and OLAP in a single database.

MySQL HeatWave is an advanced database cloud service with significant additional capabilities and benefits that set it apart from other open source and commercial cloud database offerings, making it much more than just another parallel database.

Automation with *Autopilot*



First, the addition of *MySQL Autopilot* to *MySQL HeatWave* provides built-in machine learning based automation for several key aspects of system operation and administration, benefitting from built-in intelligence.

MySQL Autopilot learns from the behavior of its users to enhance its own operations over time. These learnings are used to place data automatically for better performance (auto placement); to determine the best size for the cluster (auto provisioning); to build better query plans over time, as the system learns the distribution of data values in the database (auto query plan improvement); and, to prioritize short queries while still ensuring that longer queries get good service (auto scheduling). Other operations, such as data loading and system error recovery, are also automated to further simplify the building and operation of databases.

Scale



Previous cloud services based on MySQL offer some limited forms of parallel or clustered database operation. By contrast, *MySQL HeatWave* is a fully distributed, massively parallel architecture, capable of processing queries in parallel across 64 nodes (servers). Further, in *HeatWave* Oracle has designed and implemented new algorithms for distributed query processing. This means that queries over a large amount of data can be executed much more rapidly (as much as 5400 times faster) than with other implementations of MySQL,

About WinterCorp

WinterCorp is an independent consulting firm, expert in the strategy, architecture and scalability of the modern analytic data ecosystem.

Since our founding in 1992, we have architected and engineered solutions to some of the toughest and most demanding analytic data challenges, worldwide.

We help customers define their data-related business interests; develop their data strategies and architectures; select their data platforms; and, engineer their solutions to optimize business value.

Our customers get business results with analytics in which their return is often ten or more times their investment.

When needed, we create and conduct benchmarks, proofs-of-concept, pilot programs and system engineering studies that help our clients manage profound technical risks, control costs and reach business goals.

With our in-depth knowledge and experience, we deliver unmatched insight into the issues that impede scalability and into the technologies and practices that enable business success.



WinterCorp

www.wintercorp.com

**TYNGSBORO, MA
617-695-1800**

and at a fraction of the cost. It also means that the throughput of the total system is much higher than can be achieved on a single server or small cluster of servers.

In addition, *MySQL HeatWave* supports automatic replication of data from an on-premises source system running on another copy of MySQL, a far more efficient and streamlined process than that offered with MySQL on other clouds.

Performance and Price/Performance

Oracle has published benchmark data showing that *MySQL HeatWave* delivers much higher performance and much better price performance than comparable cloud database services offered by other leading cloud database vendors, such as Amazon, Google, and Snowflake.

While WinterCorp has not independently validated these benchmarks, they do seem to indicate very large advantages in both performance and cost of operation. Even the best, independently audited, standard benchmarks are only suggestive of what customers might experience with their own data and queries, so customers should always interpret standard benchmarks with caution.

Nevertheless, with the fully transparent publication of these tests, Oracle has quantified its claims for performance and price performance and provided publicly available scripts and evidence to support them. Customers considering *MySQL HeatWave* should perform their own evaluations and tests before relying on any particular level of performance or cost savings for their applications and workloads.

Recommendation

CUSTOMERS WITH AN INTEREST in open-source database cloud solutions, including data warehouse solutions, ought to take a close look at *MySQL HeatWave*. This is MySQL in a way we have never seen it before, with highly parallel query execution, significant machine learning-based automation with *MySQL Autopilot*, demonstrated excellent price performance compared to other offerings, and other strengths. Here customers gain the benefits of open source with vendor backed support integrated into Oracle Cloud Infrastructure built for high security, high performance database operations. On the basis of the *MySQL HeatWave* architecture alone, no other MySQL option currently available comes close in query performance.

I don't recommend *MySQL HeatWave* for enterprise-scale or mission-critical data warehouse requirements for which one would currently use Oracle Database. However, *MySQL HeatWave* is a clearly superior solution for open-source database users looking for cost-effective, high performance for their application-driven data warehouses while avoiding the headaches of managing separate databases for transaction processing and analytics. ●