

OFFERING OVERVIEW

Oracle Switches to Autopilot, Turns Up the Heat With MySQL Database Service

Oracle Adds Machine Learning Automation to
MySQL HeatWave



Holger Mueller
Vice President and Principal Analyst

TABLE OF CONTENTS

Executive Summary..... 3

About Oracle MySQL Database Service With HeatWave 4

Functional Capabilities 9

Analysis and Observations17

Recommendations.....19

Related Research20

Endnotes23

Analyst Bio24

About Constellation Research.....25



EXECUTIVE SUMMARY

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

Business Themes



Data to Decisions



Technology
Optimization



New C-Suite

ABOUT ORACLE MYSQL DATABASE SERVICE WITH HEATWAVE

Overview

MySQL became part of Oracle via the Sun Microsystems acquisition in 2009.¹ 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's latest MySQL innovation was the December 2020 addition of HeatWave, a cloud-native in-memory query accelerator designed to accelerate OLAP and other complex queries within Oracle MySQL Database.

As a further sign of Oracle's continued investment and stewardship of MySQL, Oracle is adding machine learning-based automation to HeatWave, called MySQL Autopilot, bringing automation capabilities to an open-source cloud-native database without any additional charge. The most prominent features of MySQL Autopilot 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 us that database choices matter. Why? Switching databases is hard and expensive, and doing so often poses a risk to business continuity and security. Leaders

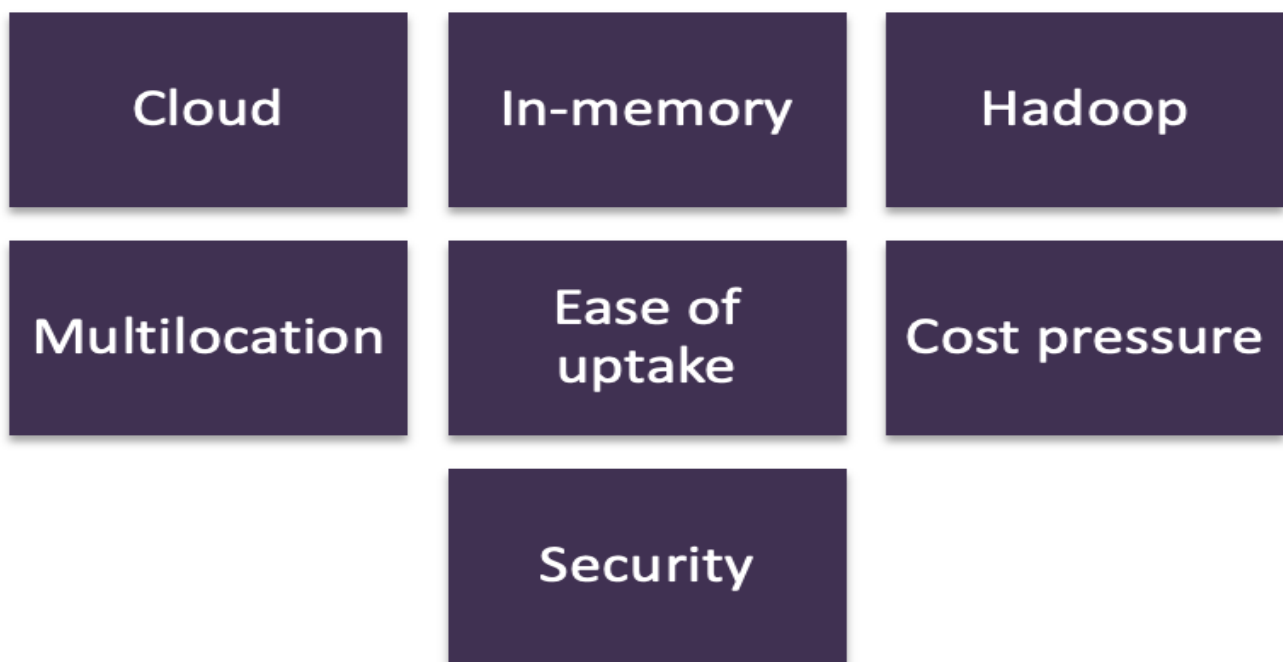
must make wise decisions about databases that provide a long-term return on investment, reduce overall operating costs, and deliver on enterprise agility.

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 databases available as on-premises products. Notably, the Hadoop-centric database vendors have been the slowest to offer public cloud deployments, relative to the start date of their first commercial offerings. CxOs expect elasticity of cloud resources, consumption-based pricing, and flexibility of licensing models.

Figure 1. Seven Market Trends Defining Next-Gen Databases



Source: Constellation Research

- 2. New designs support in-memory computing.** Moore's Law has not only made hard disk drives (HDDs) cheaper but also lowered the cost of random-access memory (RAM). The result: Enterprises can run large parts or, in some cases, the complete transactional enterprise resource planning (ERP) databases of large multinational enterprises in memory. Performance benefits include the ability to process information in real time. Traditional vendors keep a subset of data in (expensive) memory. Consequently, how memory is managed, accessed, and consumed is where the vendors differ. Hadoop-based vendors use Apache Spark and Hive. Traditional RDBMS vendors use memory sparingly and in an organic fashion, moving data to an in-memory database when beneficial for system performance or placed in memory by system administrators. In-memory vendors take a more radical approach and place the entire system in memory at all times. Meanwhile, cloud-based vendors, given the novelty of their offering, have not shared their in-memory philosophy (Oracle being an exception with their HeatWave announcement); however, given the use cases, Constellation expects the uptake of in-memory options to be similar to that of traditional RDBMS vendors.

- 3. Options must coexist and integrate with Hadoop.** Hadoop has profoundly changed the data storage and retrieval market while massively transforming enterprise best practices for analytics. For the first time, enterprises can store all of their electronic information in one place, without knowing what they want to analyze and 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 machine learning (ML) and artificial intelligence (AI) will improve implementation success.


- 6. Open-source options increase cost pressures on everyone else.** With the rise of mostly open-source, cloud-based databases, the cost pressure on the other vendors has increased. Vendors see themselves forced to justify the license costs of their traditional databases when the base version of an open-source-based database is free. And the revenue potential and wallet share for the add-ons that can be licensed in the open-source market are not the same as for traditionally expensive database licenses. All traditional database vendors can already see the pressure on their license prices. Constellation expects this trend to continue.
- 7. Security remains paramount to clients.** Enterprise-grade use requires a set of security capabilities to prevent data breaches, cyberattacks, and ransomware to avoid potential liability challenges for an enterprise. But enterprises are also ready to experiment and evaluate when only limited security concepts are available; Hadoop databases are an example. In the end, the value of insight gleaned from data often trumps security. However, large-scale and production deliveries require an acceptable set of security capabilities. To an increasing degree, enterprises expect security capabilities to be automatically deployed. Finding the right balance between insight and security is key for CxOs.

The Machine Learning Opportunity

ML (also known by the alternate marketing term AI) is revolutionizing software. Powered by the trend toward Infinite Computing,² ML enables for the first time the ability of software to act for humans, and with that taps into a tremendous automation potential.

Traditionally, automation models, often powered by predictive analytics, would be a labor of love for data scientists. By using long phases of data modeling, hypothesis testing, model validation, and testing, data scientists provided automation that was model-based. The challenge for today's enterprises is twofold: On the one hand, there is not the time to wait for multiple months to create working automation; on the other hand, there are not nearly enough data scientists to provide all the ML enterprises need to succeed in the digital economy.

Figure 2. Eight Market Trends Define the EAP Market



Source: Oracle

The development to address this need has been the creation of software architectures and frameworks that allow for the creation, validation, and operation of ML models. It is “automatic” ML—in short, AutoML, as Oracle calls it. Practically, AutoML provides a human-free pipeline that accelerates the creation and adoption of ML models for specific problems. Interestingly, Oracle AutoML was developed by the same team that developed MySQL HeatWave.

Oracle faces the challenge that it wants to offer a MySQL cloud service for the core MySQL clientele, which is not ready to afford big spending on the operationalization of its databases in the cloud. Consequentially, Oracle has used AutoML capabilities to build machine learning–based automation capabilities for its MySQL HeatWave service called MySQL Autopilot (see Figure 2).

As a result of Autopilot, Oracle can offer MySQL on OCI in a very cost-efficient way, with little to zero human operational involvement. It is almost as if Oracle took a page from the playbook of MySQL’s bigger “brother” Oracle Database but with more emphasis on automation that is driven by machine learning. Using machine learning to database automation (Autopilot) was greatly facilitated by Oracle AutoML, which was designed in parallel by the same team at Oracle Labs.

FUNCTIONAL CAPABILITIES

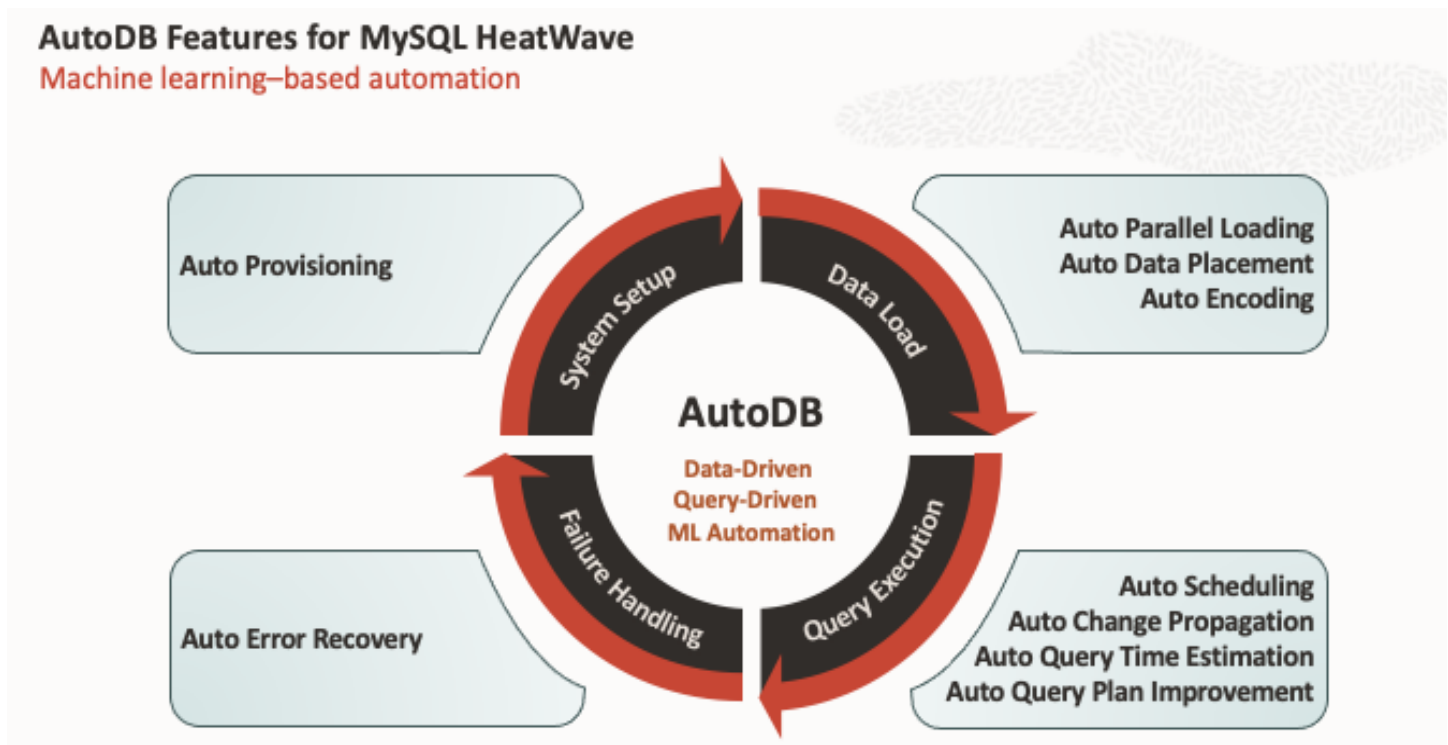
Best of Both OLTP and OLAP in MySQL Database Service With HeatWave, Now Powered by Autopilot

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

- **Preserve standard SQL as access language.** Standard (unaltered) SQL supported by MySQL had to remain the common programming language for both transactional and analytical workloads.
- **Keep it 100 percent backward-compatible.** No code changes could be required for running existing MySQL workloads with the new architecture. This would ensure that the vibrant third-party ecosystem of ISVs could operate their offerings seamlessly, with no code changes or testing.
- **Achieve extreme performance.** Performance of analytic workloads while not compromising transactional performance was critical.
- **Optimize for a commodity cloud.** The system had to be designed for the cloud—to achieve cloud scalability, use commodity cloud services (interconnect, storage, compute), and obtain best performance at the lowest cost.
- **Focus on machine learning-based automation.** Workload-driven ML was necessary to intelligently automate operations such as provisioning, optimization of query execution, and other services.

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

Figure 3. The Autopilot Features of MySQL HeatWave



Source: Oracle

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

Equally important was Oracle's ability to make the HeatWave engine scale really well. Using a massively parallel architecture, HeatWave achieved the design goal of 5,400x faster performance than with standard MySQL on 4TB of data that could be running on-premises or in any cloud. As the data size increases, the performance difference increases beyond 5,400x.

Oracle implemented Autopilot in HeatWave for the following key database events (see Figure 3):

- System setup is facilitated by Auto Provisioning
- Data load is automated by Auto Parallel Loading, Auto Data Placement, and Auto Encoding

- Query execution is powered by Auto Scheduling, Auto Change Propagation, Auto Query Time Estimation, and Auto Query Plan Improvement
- Failure handling is improved by Auto Error Recovery

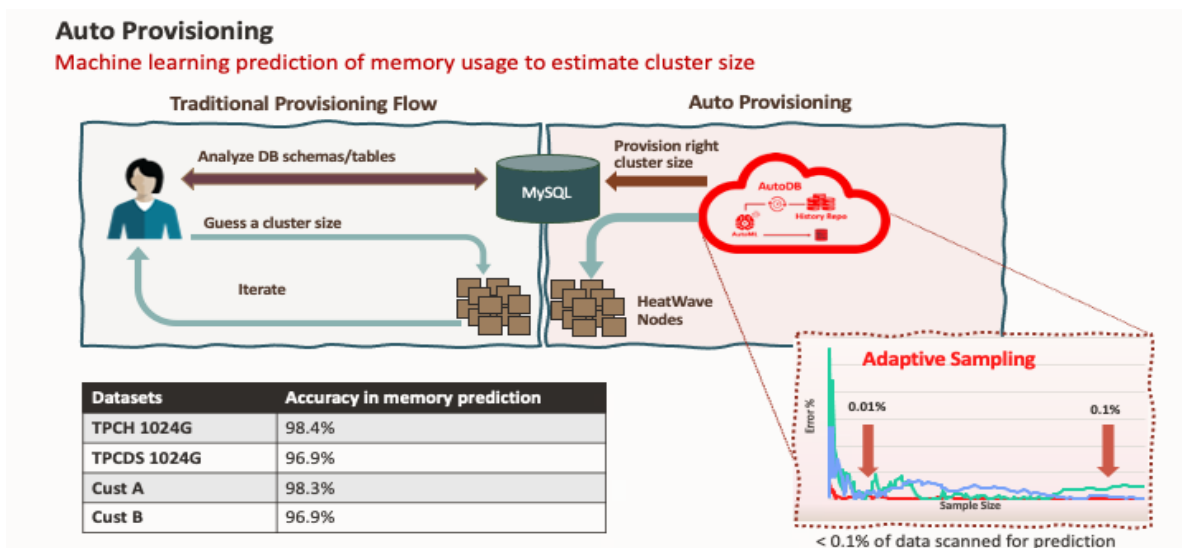
Auto Provisioning Determines the Optimal Size of the Cluster Automatically

Allotting hardware resources has been tricky throughout the history of databases. Before the cloud, system engineers took care of this effort by sizing the underlying hardware for the expected usage of the database, and database administrators worked within those confines and manually tuned the database to meet performance goals as workloads changed. Configurations were static during that time, because changing them often meant downtime that was always related to operational risk.

Obviously, this is not a workable approach in the cloud era, where cloud infrastructure needs to support thousands of clients' MySQL instances, with very different workloads and constantly changing load profiles. Consequently, it does not come as a surprise that Oracle has been looking at using machine learning to address this operational challenge.

With the new capability of Auto Provisioning (see Figure 4) as one of MySQL Autopilot's features, an intelligent adaptive sampling mechanism predicts the amount of memory that will be required to process

Figure 4. How Oracle Enables Auto Provisioning for MySQL HeatWave



Source: Oracle

these tables—both for storing the encoded representation of the table in memory and memory needed to process queries. Based on this memory prediction, the system makes a recommendation for the size of the cluster that is required.

When Oracle benchmarked the accuracy of its new machine learning–powered capabilities provided with Auto Provisioning, it consistently saw system-managed provisioning accuracy of more than 95 percent and therefore added this very useful machine learning–powered capability to Autopilot.

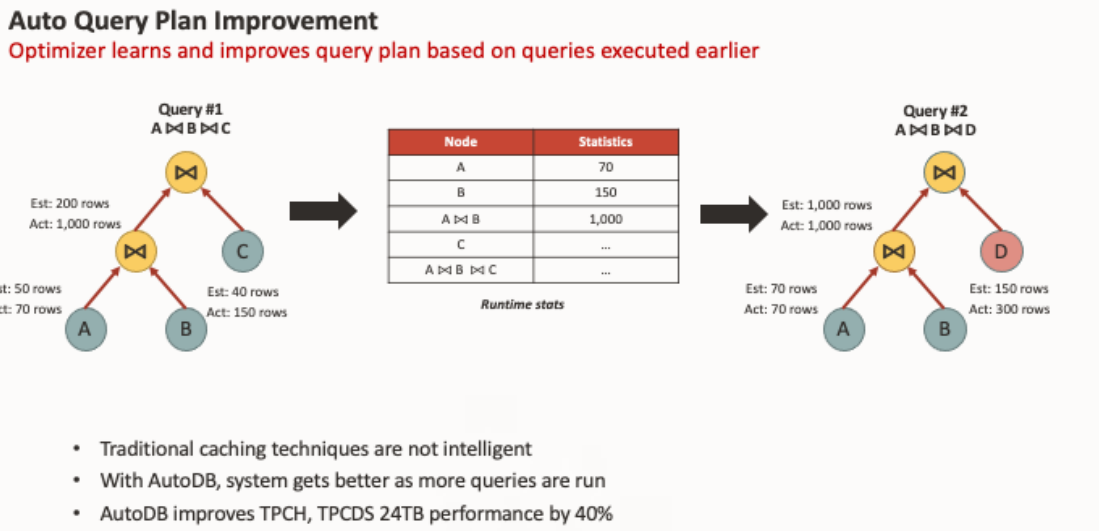
The successful operation of Auto Provisioning gives CxOs the peace of mind that the system can intelligently recommend the optimal size of the cluster. This is a capability that is not offered by other cloud services.

Auto Query Plan Improvement Boosts MySQL Performance

Database vendors have for the longest time worked on optimizing queries that users unleash on their database offerings. It is clear today that optimizing queries in combination with storage is the key to better database performance.

In the past, query optimization was hard coded. Then it became rules-based, with better options for updating the underlying, resulting query plans that tell a database how to get to the requested data

Figure 5. How Oracle Powers Auto Query Plan Improvement for MySQL HeatWave



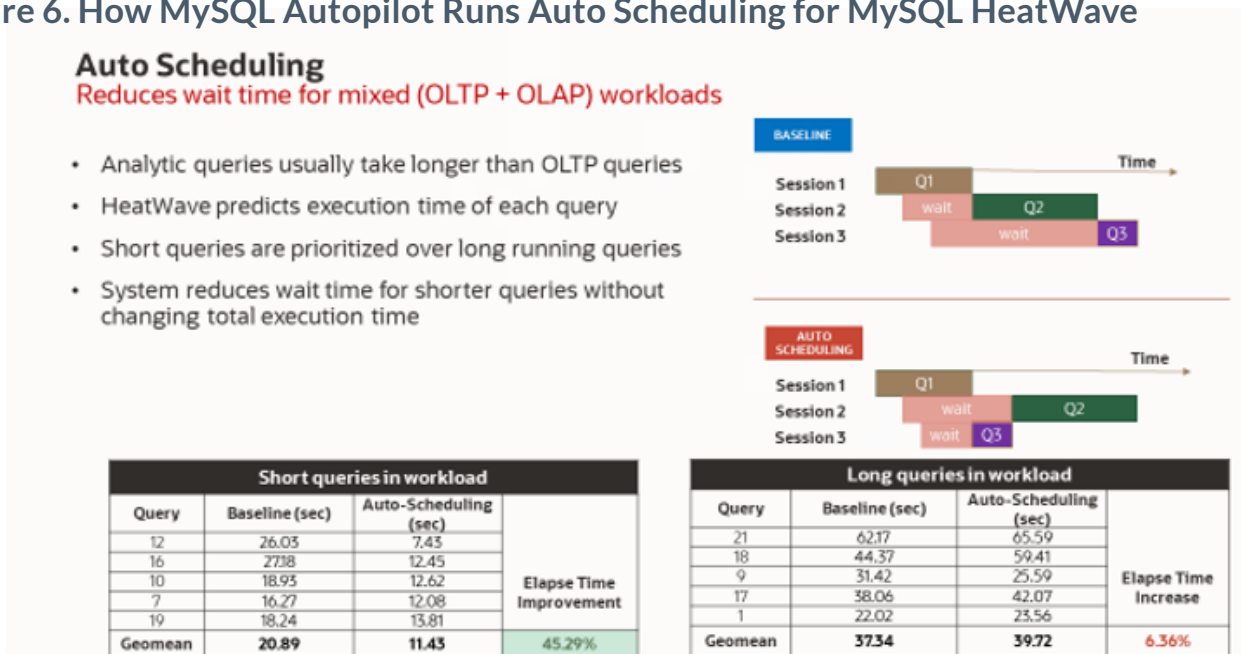
Source: Oracle

in the most efficient way. Today, the most efficient way to optimize query plans is via ML. This allows a more individualized and real-time optimization of queries, because the query optimizer learns and adapts during the database usage time. Effectively, query optimization moves from being a static construct to being a dynamic solution that is optimized for each individual database and its current, temporary usage load.

When adding its Autopilot capabilities to MySQL HeatWave, Oracle did not balk at applying ML, and the solution is Auto Query Plan Improvement (see Figure 5). Obviously, Auto Query Plan Improvement learns as more queries are being executed, taking advantage of the intelligent capabilities ML brings to query plan optimization. The user benefit for this intelligence is that, unless cache-based mechanisms which improve performance of identical queries have been cached, Autopilot improves performance of queries that it has never seen before.

The benefit of MySQL HeatWave’s Auto Query Plan Improvement is proven: In the well-known performance tests TPCH and TPCDS, it improved performance on a 24TB system by 40 percent. Features such as these are key for CxOs, because they give them the peace of mind to entrust their next-generation applications to MySQL HeatWave, with the end goal in mind: Enterprise Acceleration without additional cost.³

Figure 6. How MySQL Autopilot Runs Auto Scheduling for MySQL HeatWave



Source: Oracle

Auto Scheduling Automates MySQL to New Levels

Successful databases get bombarded by queries left and right. For MySQL HeatWave, the situation is even a little more complex, because it needs to perform both OLTP and OLAP queries, with OLAP queries typically requiring more database resources than OLTP queries.

This creates the conundrum of optimizing between OLTP and OLAP queries. Favoring the faster and typically smaller OLTP queries results in bad performance for OLAP users, who must wait longer and longer. Taking the opposite approach, favoring OLAP queries means that critical transaction loads on the OLTP side may take too long, and even time out in the worst cases.

MySQL HeatWave predicts query times for both OLTP and OLAP queries, and then uses (no surprise by now) ML to schedule these queries via the new Autopilot feature, Auto Scheduling (see Figure 6).

By prioritizing queries with lower execution time, Oracle has managed to optimize overall system responsiveness. The result is a net gain in overall performance for MySQL HeatWave, which should overall make MySQL HeatWave users happy with using the database. OLAP users may wait a little longer (reports can usually be a few minutes late), knowing they will not be blamed for (too) slow OLTP transactions and, in the worst case, even transaction time-outs. On the other side, the transaction performance gains outweigh the slowdown on analytical queries, as Figure 7 illustrates.

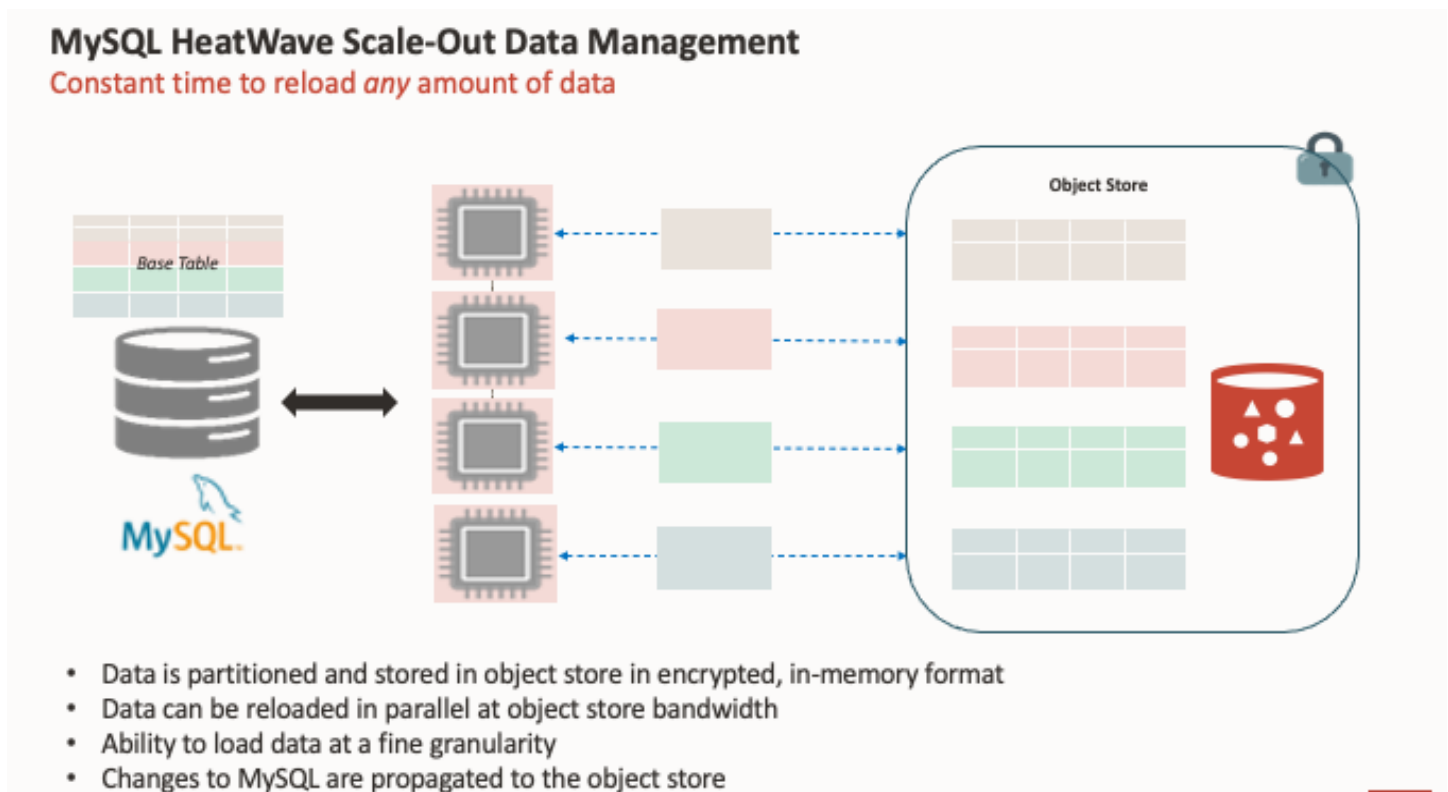
How Oracle Achieves Leading Price/Performance for MySQL HeatWave

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

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

- 1. Cost efficiencies by leveraging OCI.** HeatWave uses the most cost-effective commodity services from OCI, including object storage, commodity interconnect, and flexible compute shapes, and benefits

Figure 7. MySQL HeatWave's Scale-Out Data Management Architecture



Source: Oracle

from OCI's high availability, elasticity, and security as well as from the economies of scale. It's priced less than all comparable cloud database services, so talk about value for your money—HeatWave is off the charts.

2. **Use of cost-efficient computing elements.** HeatWave uses OCI shapes that provide the lowest cost per TB of memory and Oracle has optimized the software stack to achieve a great balance of compute, memory, and network processing. HeatWave also leverages the flexible VM shapes offered by OCI.
3. **In-memory hybrid columnar processing for maximum performance.** HeatWave uses hybrid columnar processing to provide the performance that analytical and mixed workloads demand. It achieves the performance via vectorized execution of the relevant columns, which can be broken into different vectors and multiple chunks as needed. HeatWave manages the vectors in chunks, which benefit from multicore scalability inside of a single partition. Beyond that, HeatWave scales horizontally across n partitions.

4. **A massively parallelized architecture lifts performance to unheard-of levels.** HeatWave achieves massive parallelism via massive partitioning, intelligently assigning workloads to multiple CPU cores, and then combining back the results. The partitioning architecture is designed for high-fanout workloads and can be done at near-memory bandwidth; the machines and CPUs used can further partition the data in parallel (see No. 3 above), and with Oracle designing the entire architecture, HeatWave can optimize partition sizes based on the cache size of the hardware.
5. **Scale-out data management boosts performance and reliability.** By partitioning data and storing it encrypted in the object stored in the HeatWave in-memory format, MySQL HeatWave gets a big performance boost and improved availability. Data is encrypted for security reasons, and since the in-memory format doesn't require any additional transformation, HeatWave can read the data at near object store bandwidth, and the data can be read from the object store in parallel by all the nodes of the HeatWave cluster. Together, this results in a tremendous speedup for the load times. A 10TB cluster, which earlier would take 7.5 hours to reload, can now be reloaded in four minutes in the fastest OCI regions (see Figure 7). Moreover, the new MySQL HeatWave scale-out data management capability allows users to load data at fine granularity, reducing overall hardware expense and improving performance. No matter where information changes are stored in memory, they get propagated back to the physical object store for persistency. The overall result is massively improved system performance and reliability.
6. **The “chip-to-click” benefits at work—algorithms distributed for OCI.** Oracle has achieved tremendous performance and TCO gains from a vertically integrated stack—from the silicon all the way through to the user's click (we refer to this as the Oracle chip-to-click stack)—and it delivers in spades for HeatWave as well. Specifically, HeatWave partitions the data to fit into the cache size, processes the partitions as fast as possible—for example, by using hand-tuned primitives such as Advanced Vector Extensions 2 (AVX2) for OCI CPUs—and optimizes the network for OCI interconnects by infusing intelligence into the scheduling of execute and transfer workloads. This helps HeatWave achieve excellent scalability across a large number of servers while using commodity hardware and services.

ANALYSIS AND OBSERVATIONS

Constellation sees the following strengths and weaknesses for Oracle's MySQL HeatWave with MySQL Autopilot offering (see Table 1):

Strengths

- **No ETL and excellent query performance to accelerate insight to action.** Enterprises need to accelerate and become more agile, and that starts with decision management. Insight to action has long eluded CxOs as MySQL cloud databases often require ETL to a separate analytical database, thereby delaying insights that are at this point based on stale data. MySQL Database Service with HeatWave makes insight to action tangible, enabled by its backward compatibility, high performance, and low cost, and with subsecond transactional changes reflected in analytical insights.
- **Ready with zero code changes.** The backward compatibility with no need to change existing OLTP code comprises an elegant architectural approach and enables an immediate uptake of the MySQL Database Service with HeatWave capabilities on top of MySQL transactional data.
- **Fully integrated with other Oracle services.** With Oracle MySQL Database Service being fully integrated with other Oracle Cloud services on OCI, the ecosystem synergies fully come into play. The full integration also eases potential concerns from CxOs toward the future investment in and viability of MySQL Database Service.
- **Autopilot is another testament to Oracle's TCO-focused organizational DNA.** With the addition of ML-based automation capabilities to MySQL HeatWave, Oracle delivers an innovation that is another proof point of Oracle's corporate DNA: It is all about lowering the TCO of running IT, in this case with MySQL HeatWave. MySQL Database Service follows in those footsteps, providing leading category performance while offering lower costs via better automation, thanks to several Autopilot capabilities added in the last release.

Weaknesses

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

Table 1. MySQL HeatWave Strengths and Weaknesses

STRENGTHS	WEAKNESSES
<ul style="list-style-type: none">• Excellent query performance to accelerate insight to action• Ready with zero code changes• Fully integrated with other Oracle Cloud services• Autopilot is another testament to Oracle's TCO-focused organizational DNA	<ul style="list-style-type: none">• A new offering• Available only in Oracle Cloud Infrastructure• It is built. How many will come?

Source: Constellation Research

RECOMMENDATIONS

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

- 1. Enable Enterprise Acceleration.** Enterprises need to move faster than ever before, and IT/computing infrastructures cannot continue to be the shackles on agility that they have been in the past. Therefore, CxOs should look at any information technology that allows their enterprises to accelerate. Not having to change code to adopt MySQL Database Service with HeatWave is a major benefit and will help the adopting enterprise to move faster and accelerate with automation provided by MySQL Autopilot.
- 2. Enable and then practice insight to action.** Enterprise decision-making has been hampered by the delay of OLTP data being available in OLAP systems. MySQL Database Service with HeatWave puts an end to this from a technology perspective. With the subsecond inclusion of transactional changes in analytical decisions, CxOs can equip their fellow CxOs and decision-makers with an analytical platform that allows real-time insight-to-action best practices.
- 3. Customers using MySQL or MySQL-compatible databases should evaluate MySQL Database Service with HeatWave now.** The immediate benefits of adopting MySQL Database Service with HeatWave, especially with the automation now available with MySQL Autopilot, are so compelling that existing Oracle MySQL customers regardless of on-premises or cloud deployments should immediately adopt the new offering. The benefits clearly outweigh the risks, and putting insight to action in reach for CxOs and practicing Enterprise Acceleration is a survival strategy for enterprises.
- 4. The addition of MySQL Autopilot to 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 were not attractive enough, Oracle has managed to make the offering even more attractive. Enterprises need to tap into the powers of ML to run their IT workloads faster and more efficiently, and the MySQL Autopilot capabilities of MySQL HeatWave have made this offering even more of a “must-evaluate” offering for both existing and potential new customers.

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," 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," May 2, 2019. <https://www.constellationr.com/research/why-c-suite-must-embrace-enterprise-acceleration>

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

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

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

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

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

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

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

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

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

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

For next-generation databases, see: Holger Mueller, "Constellation ShortList Next-Gen Databases: RDBMS for On-Premises," Constellation Research, February 12, 2021. <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 17, 2021. <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 17, 2020. <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," August 20, 2019. <https://www.constellationr.com/research/how-open-source-won-and-will-keep-winning>

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

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

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

For a Constellation ShortList on PaaS vendors, see: Holger Mueller, "Constellation ShortList PaaS Tool Suites for Next-Gen Apps," Constellation Research, February 19, 2021. <https://www.constellationr.com/research/constellation-shortlist-paas-tool-suites-next-gen-apps-3>

Also: Holger Mueller, "Constellation ShortList PaaS Suites for Next-Gen Apps," February 17, 2021. <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 on the history of MySQL, see: MySQL, Wikipedia: <https://en.wikipedia.org/wiki/MySQL>

² Mueller uses the term Infinite Computing to describe how computing resources have practically become infinite for enterprise purposes, effectively eliminating the need to size hardware resources. For more details, 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>

³ Constellation refers to Enterprise Acceleration as the need for enterprises to move faster and become more agile. For more, see: Holger Mueller, "Why the C-Suite Must Embrace Enterprise Acceleration," May 2, 2019. <https://www.constellationr.com/research/why-c-suite-must-embrace-enterprise-acceleration>

ANALYST BIO

Holger Mueller

Vice President and Principal Analyst

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

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

[🐦 @holgermu](https://twitter.com/holgermu) [📄 constellationr.com/users/holger-mueller](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.