Replace Five Services with One – Using MySQL HeatWave Natively on AWS

Mandy Pang
Senior Principal Product Manager
MySQL HeatWave

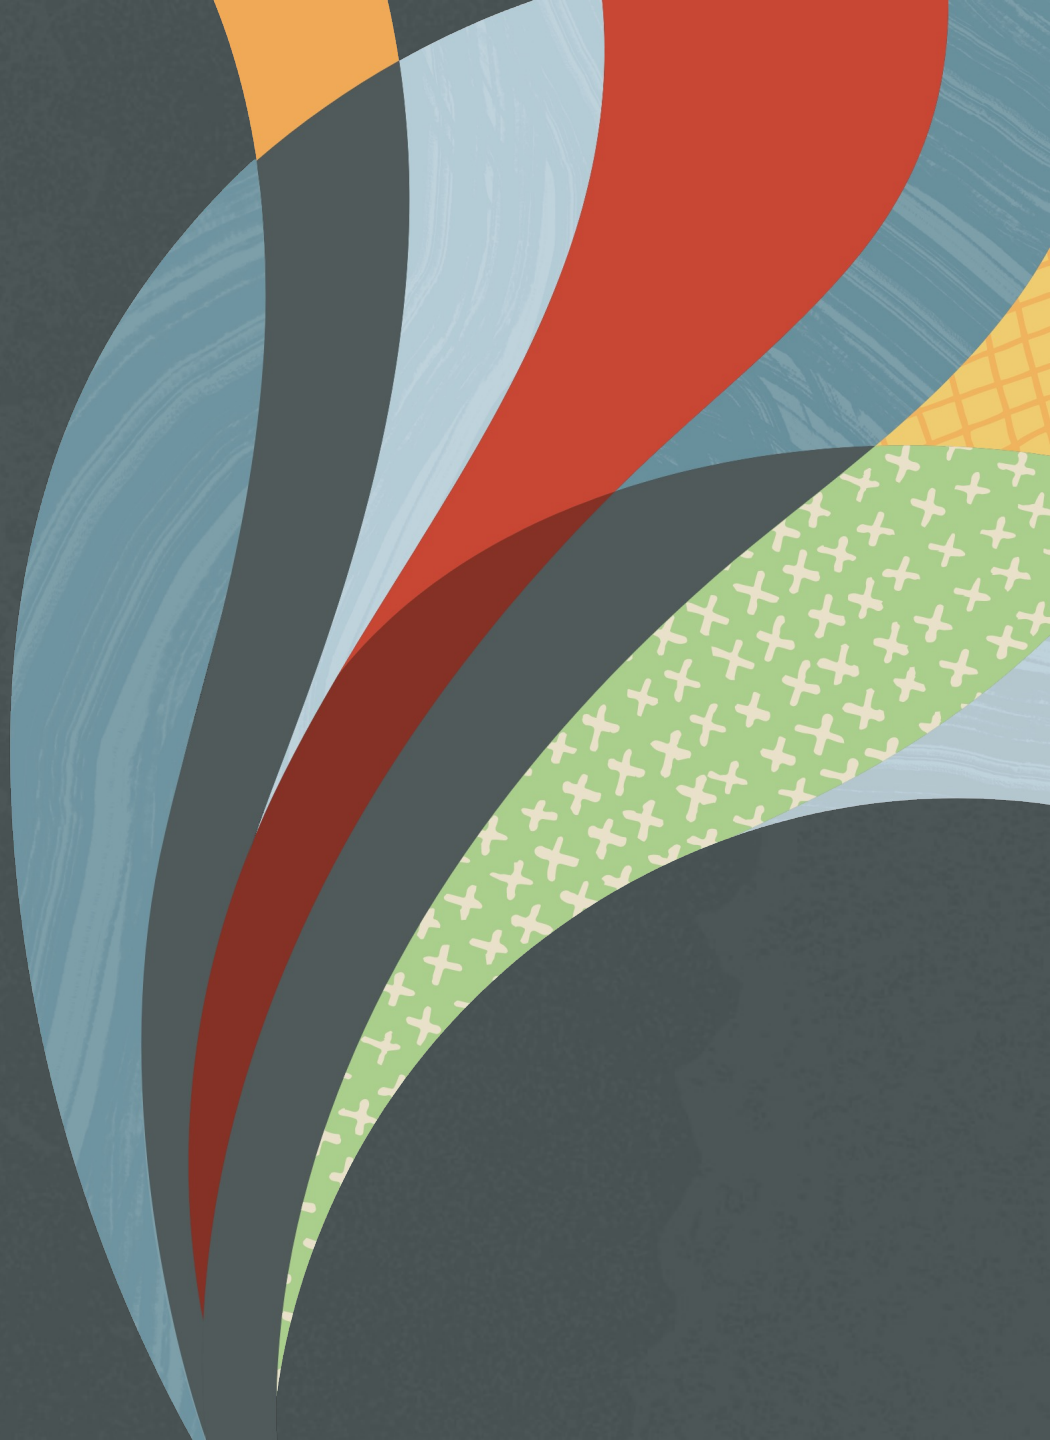

Agenda

- Complexity of using AWS services
- MySQL HeatWave on AWS
- Customer success stories

Data Platform Complexity

- **You start** with an OLTP database application

Data Platform Complexity

- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights

Data Platform Complexity

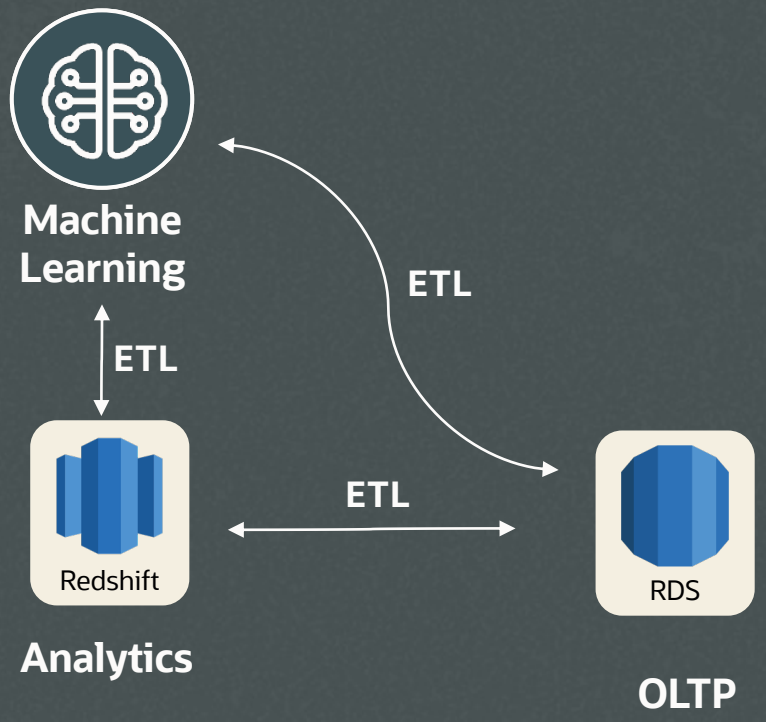

- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights
- **Machine learning** predictions will improve the customer experience

Data Platform Complexity

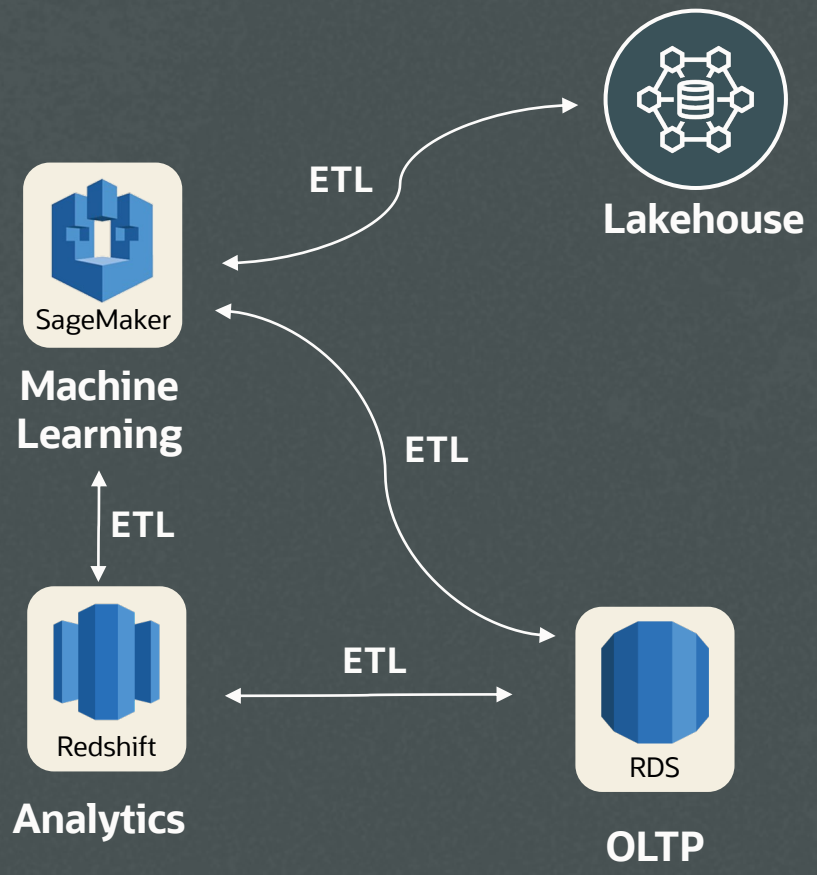

- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights
- **Machine learning** predictions will improve the customer experience
- **Lakehouse** will deliver insights into unstructured data

Data Platform Complexity

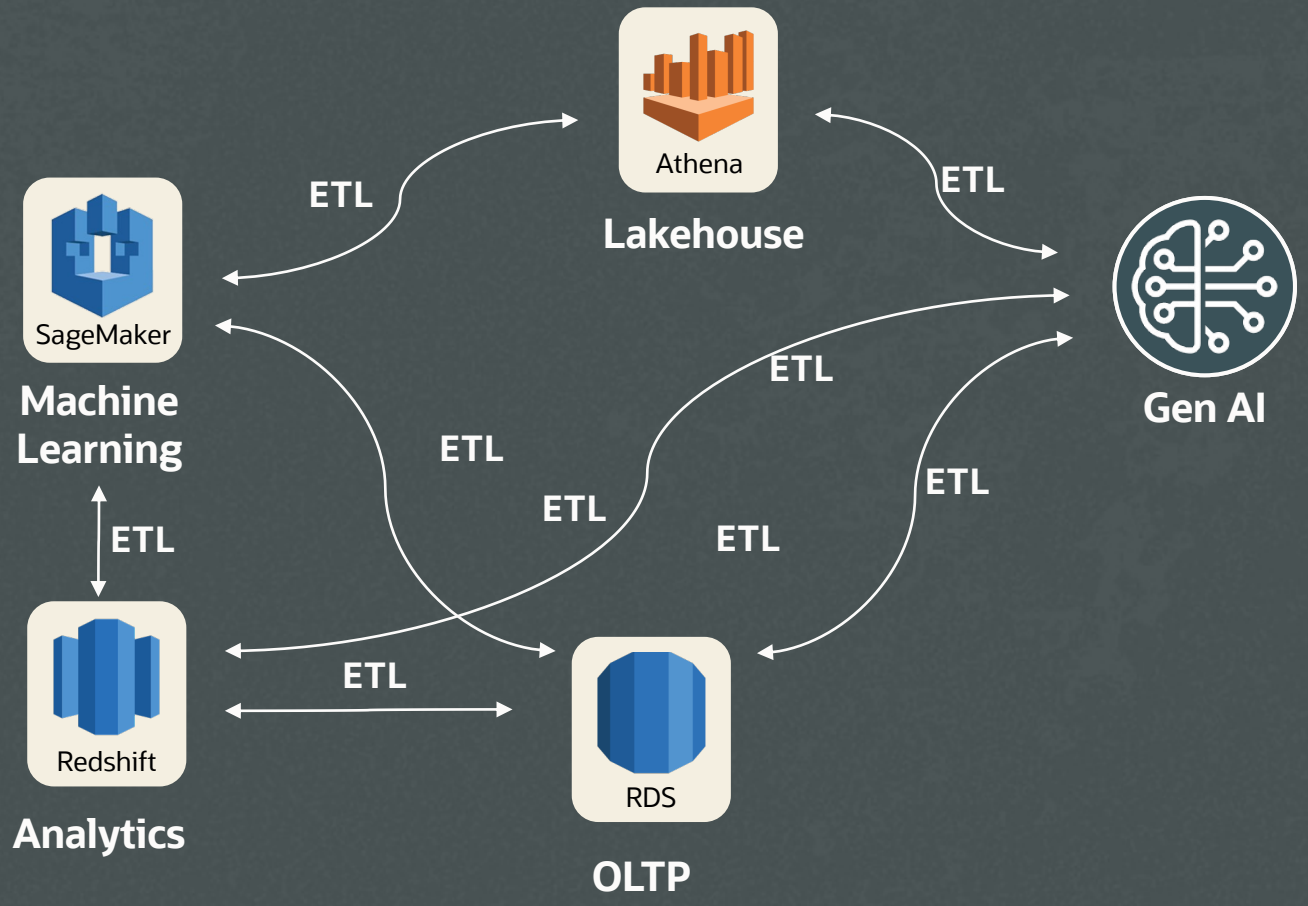

- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights
- **Machine learning** predictions will improve the customer experience
- **Lakehouse** will deliver insights into unstructured data
- **Generative AI** will give results in natural language

Data Platform Complexity

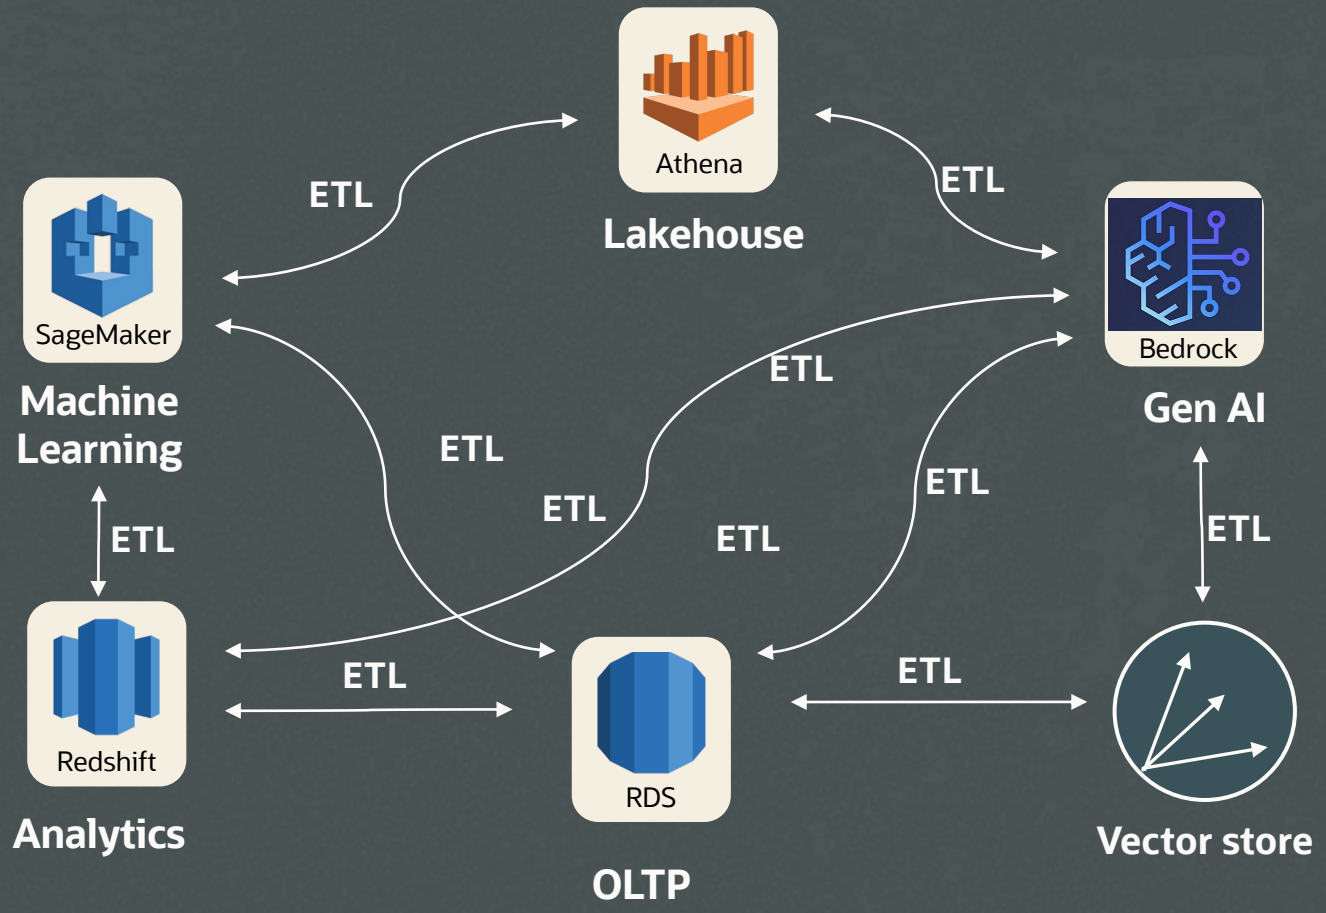

- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights
- **Machine learning** predictions will improve the customer experience
- **Lakehouse** will deliver insights into unstructured data
- **Generative AI** will give results in natural language
- **Vector store** provides context to LLM for more relevant results

Data Platform Complexity at AWS

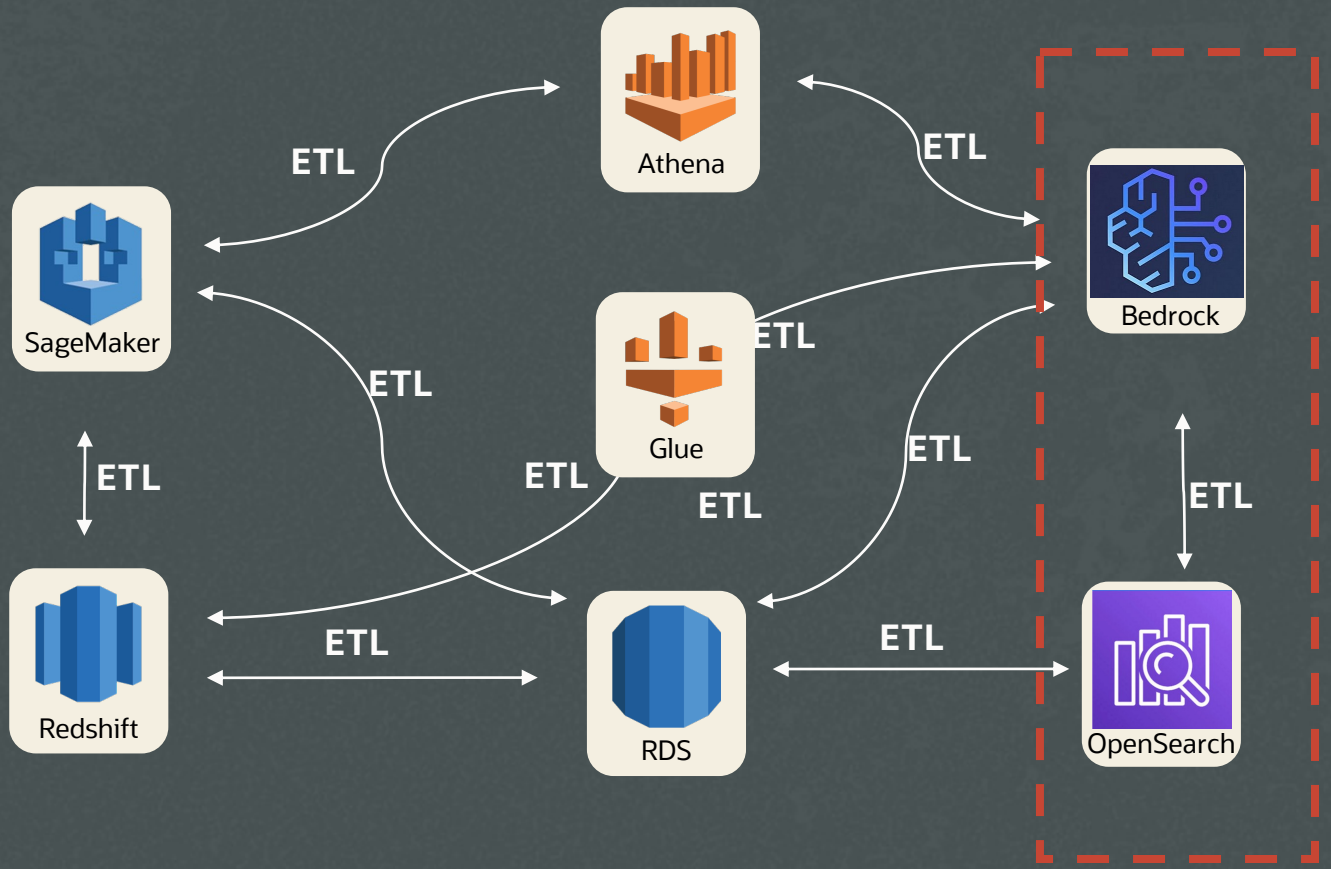

5 Separate Cloud Services

- Complex ETL processes
- Stale and obsolete data
- Difficult to maintain
- Security vulnerabilities
- Requires specialized skills

MySQL HeatWave overview

Transactions, real-time analytics, machine learning and GenAI across data warehouse and data lake in one service

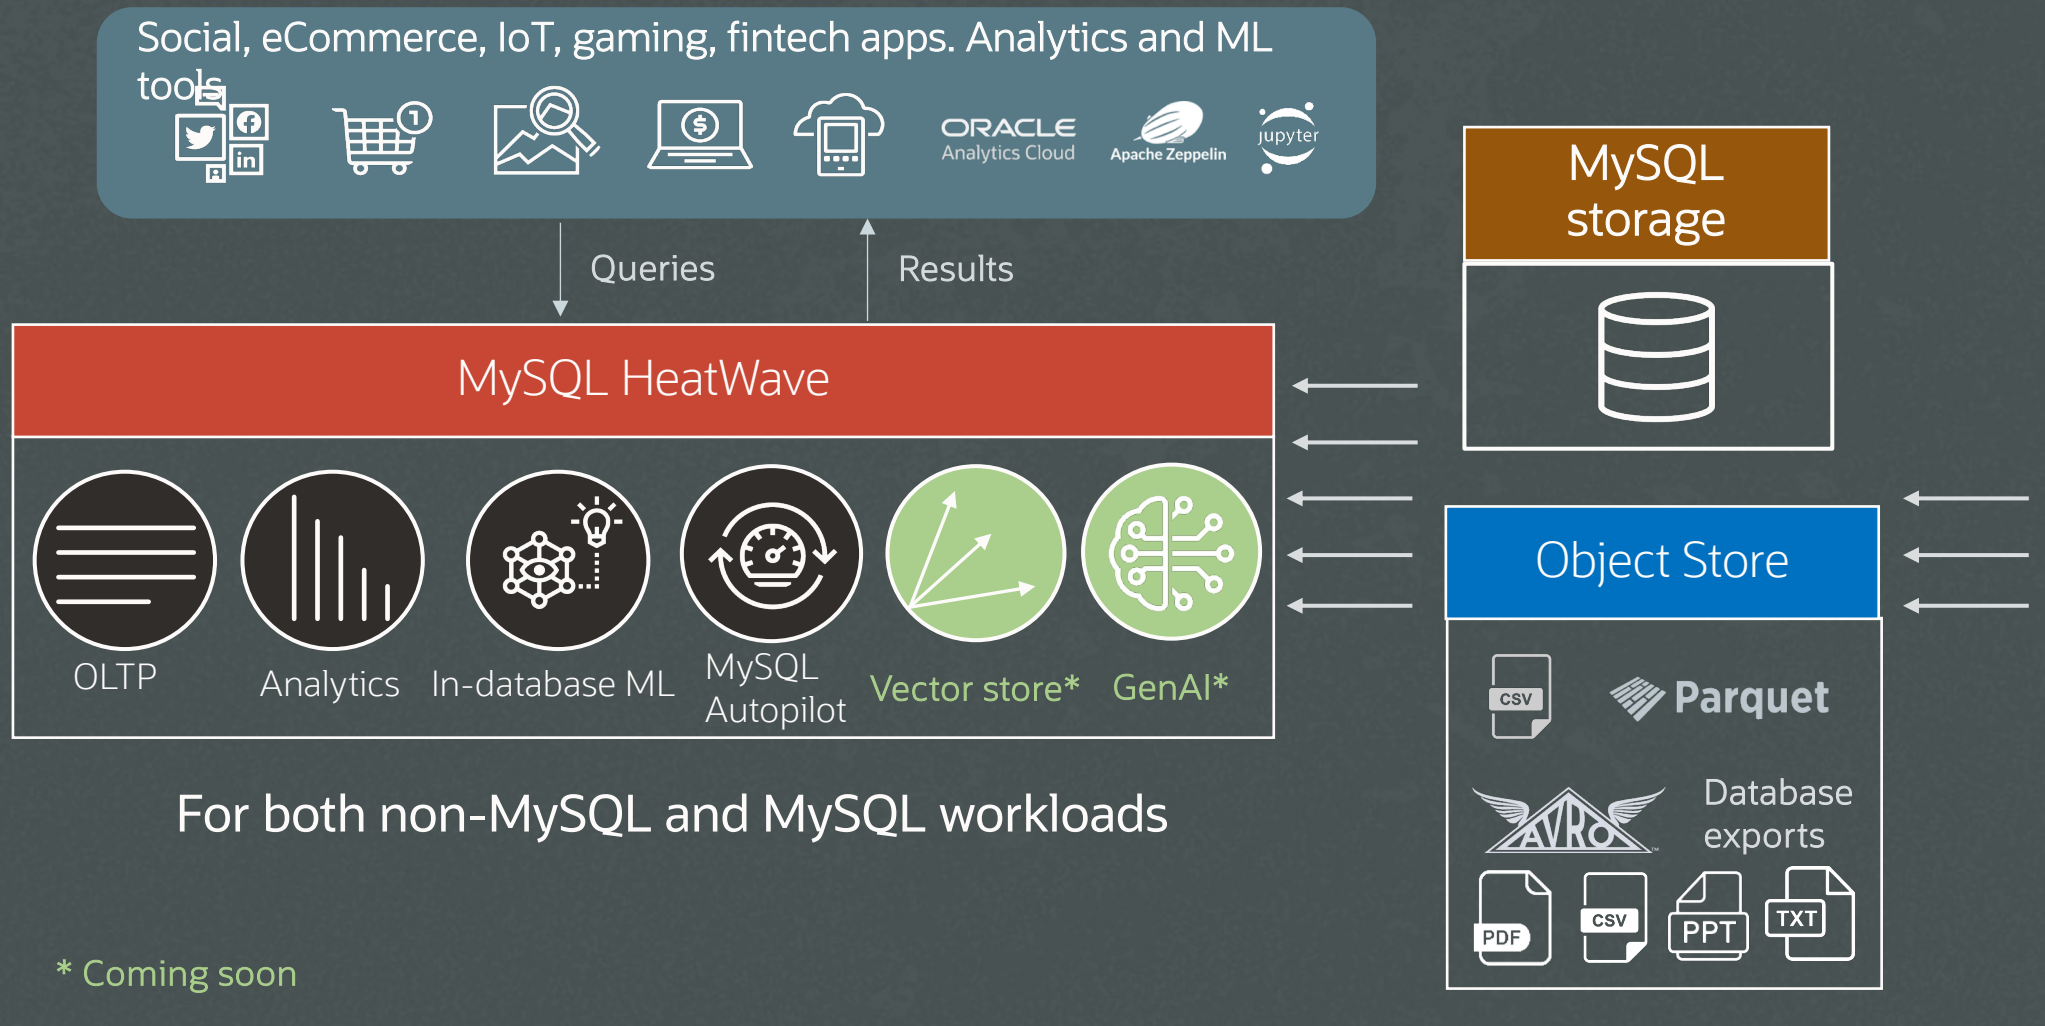

For both non-MySQL and MySQL workloads

* Coming soon

MySQL HeatWave on AWS

CONSOLE, CONTROL PLANE, AND DATA PLANE RUN NATIVELY IN AWS

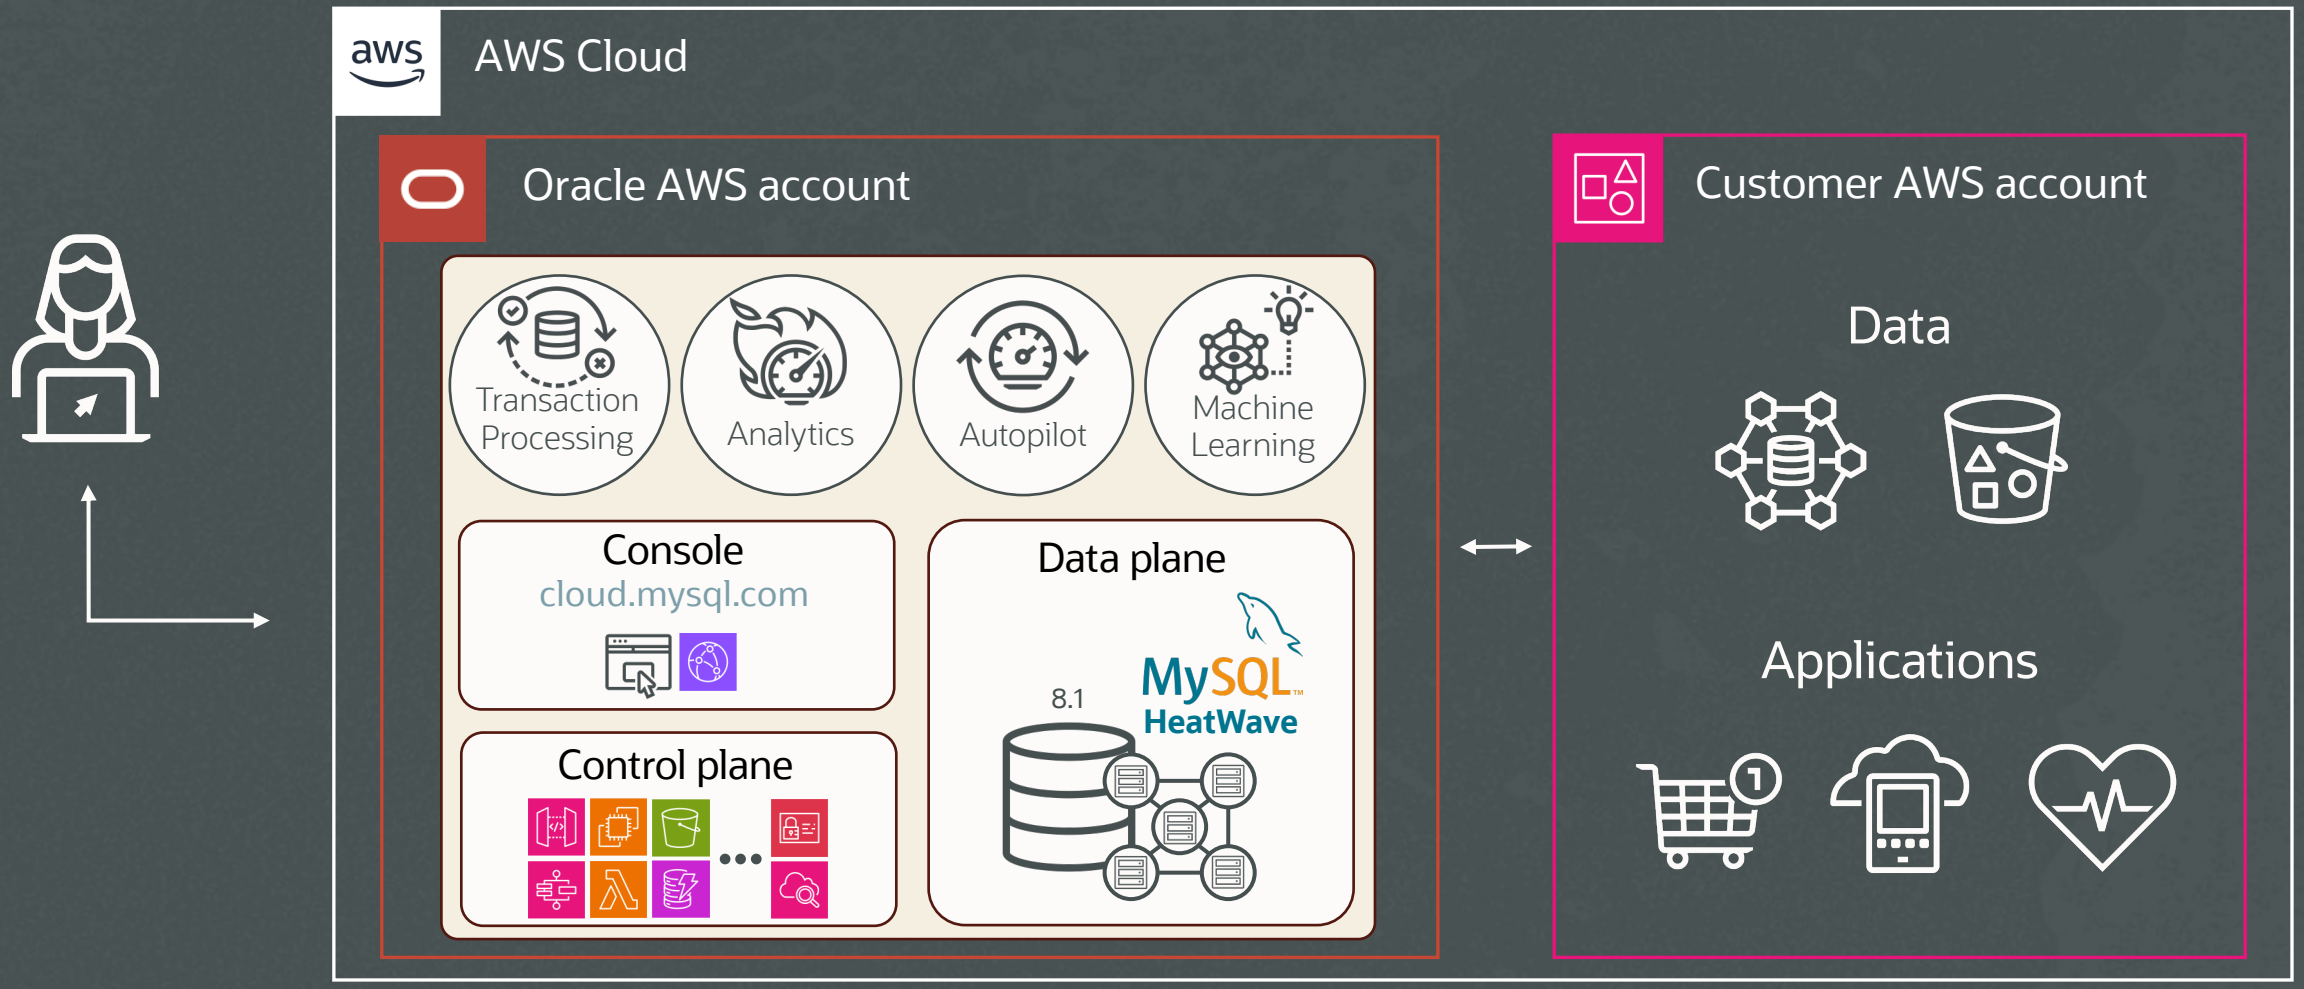

MySQL HeatWave on AWS current availability

Faster than Redshift, Snowflake, Big Query, Synapse

20X FASTER THAN REDSHIFT, 16X FASTER THAN SNOWFLAKE, 16X FASTER THAN BIG QUERY, 5X FASTER THAN SYNAPSE – TPC-H 4TB

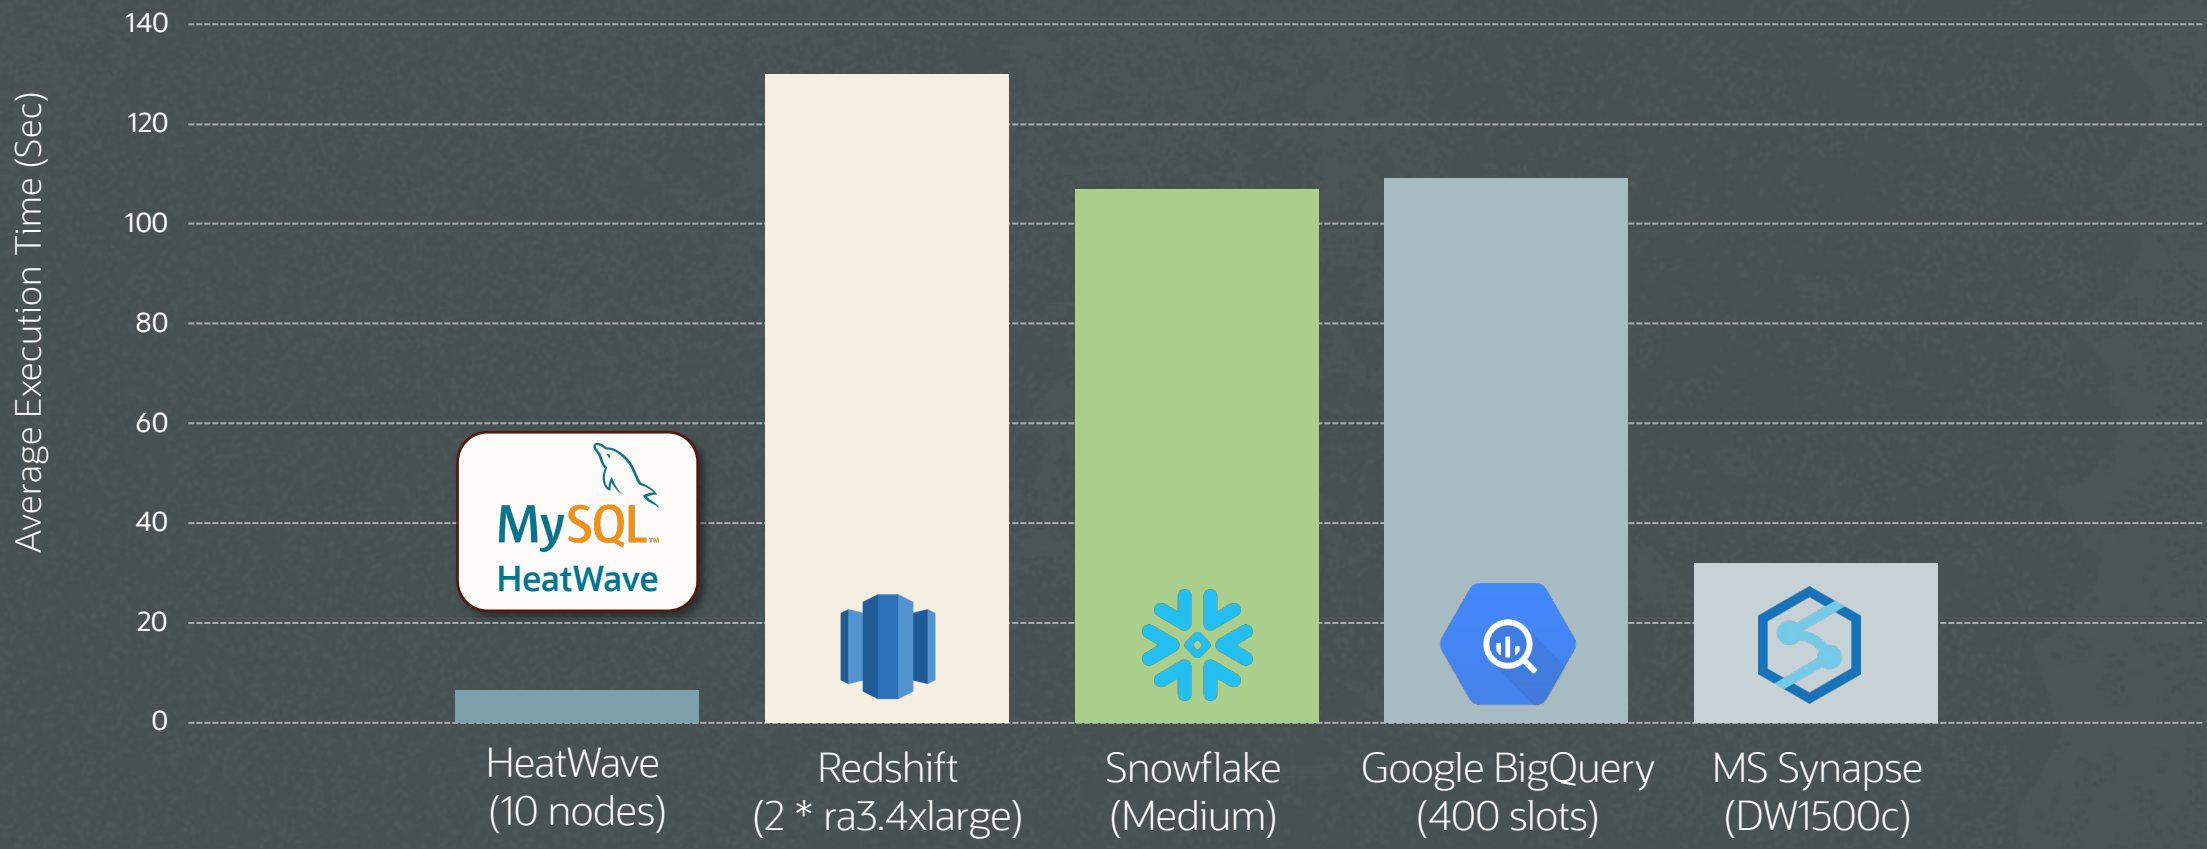

* Benchmark queries are derived from the TPC-H benchmarks, but results are not comparable to published TPC-H benchmark results since these do not comply with the TPC-H specifications.
* Results from Sept 2022

Better price performance than Redshift, Snowflake, Big Query, Synapse

7X BETTER THAN REDSHIFT, 10X BETTER THAN SNOWFLAKE, 12X BETTER THAN BIG QUERY, 4X BETTER THAN SYNAPSE – TPC-H 4TB

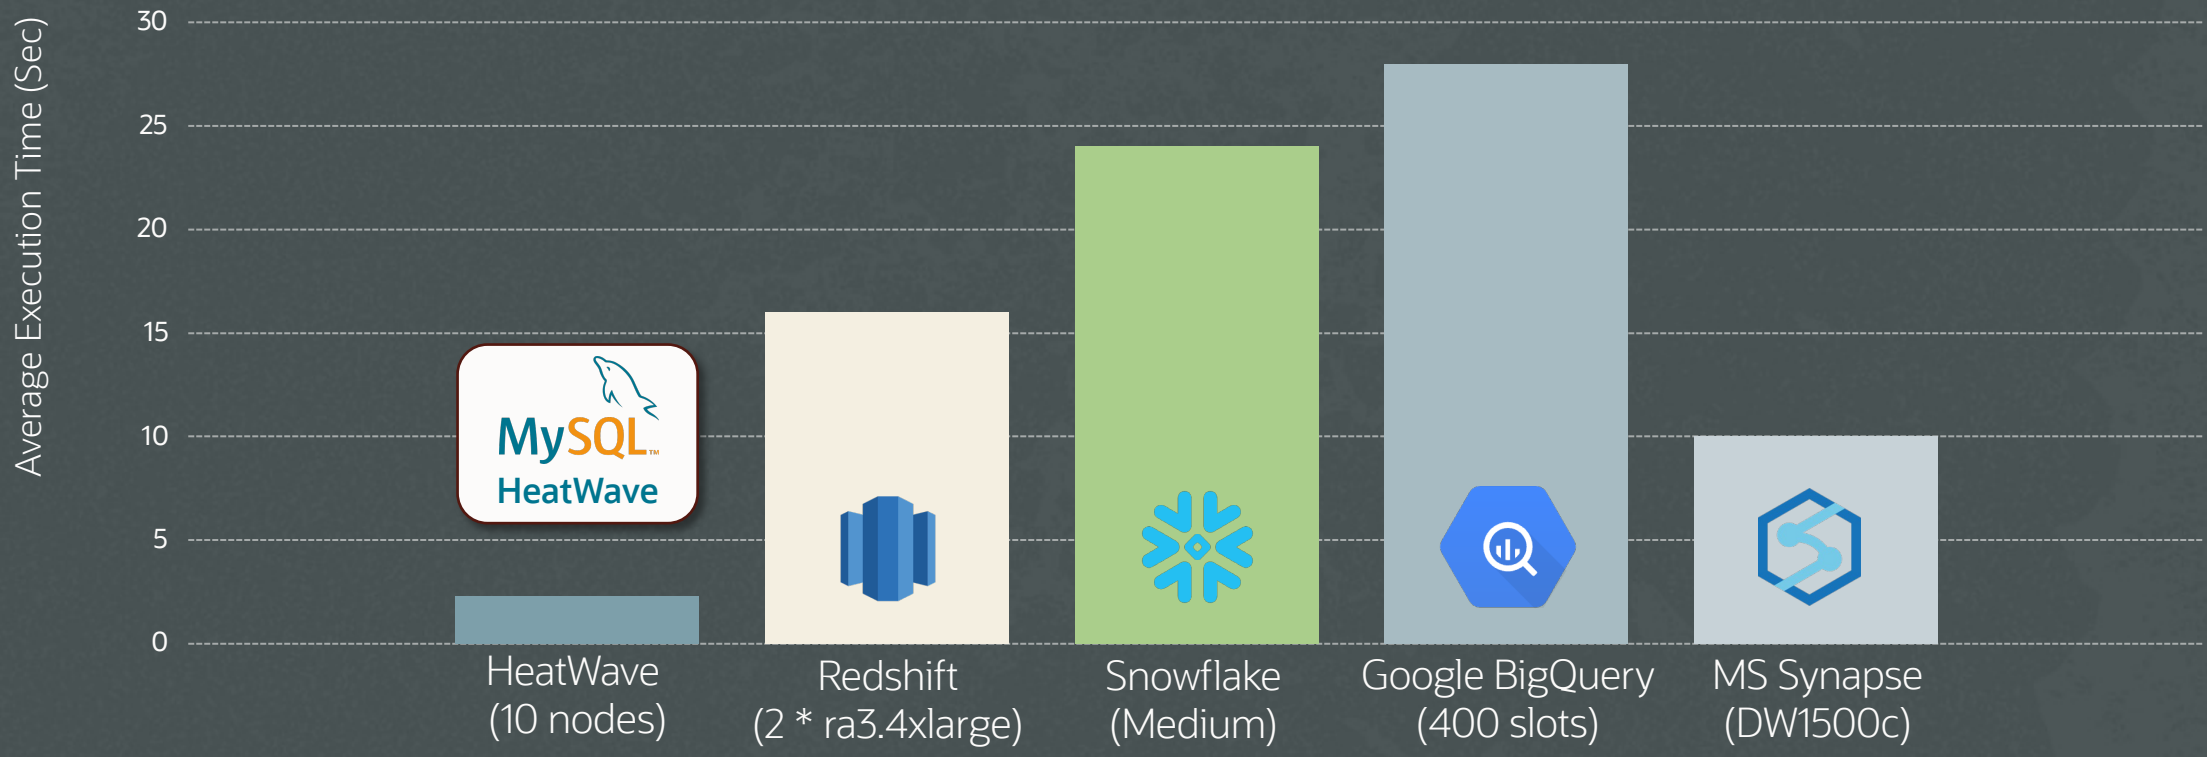

- Only compute costs are considered above
- Pricing for Redshift is based on 1-year reserved instance, paid upfront. Snowflake is based on **standard edition**
- Pricing for Google Big Query is based on monthly flat rate commitment. Azure Synapse is based on 1-year reserved pricing

* Benchmark queries are derived from the TPC-H benchmarks, but results are not comparable to published TPC-H benchmark results since these do not comply with the TPC-H specifications
* Results from Sept 2022

MySQL HeatWave offers up to 10x better throughput than Aurora for OLTP

Auto thread pooling

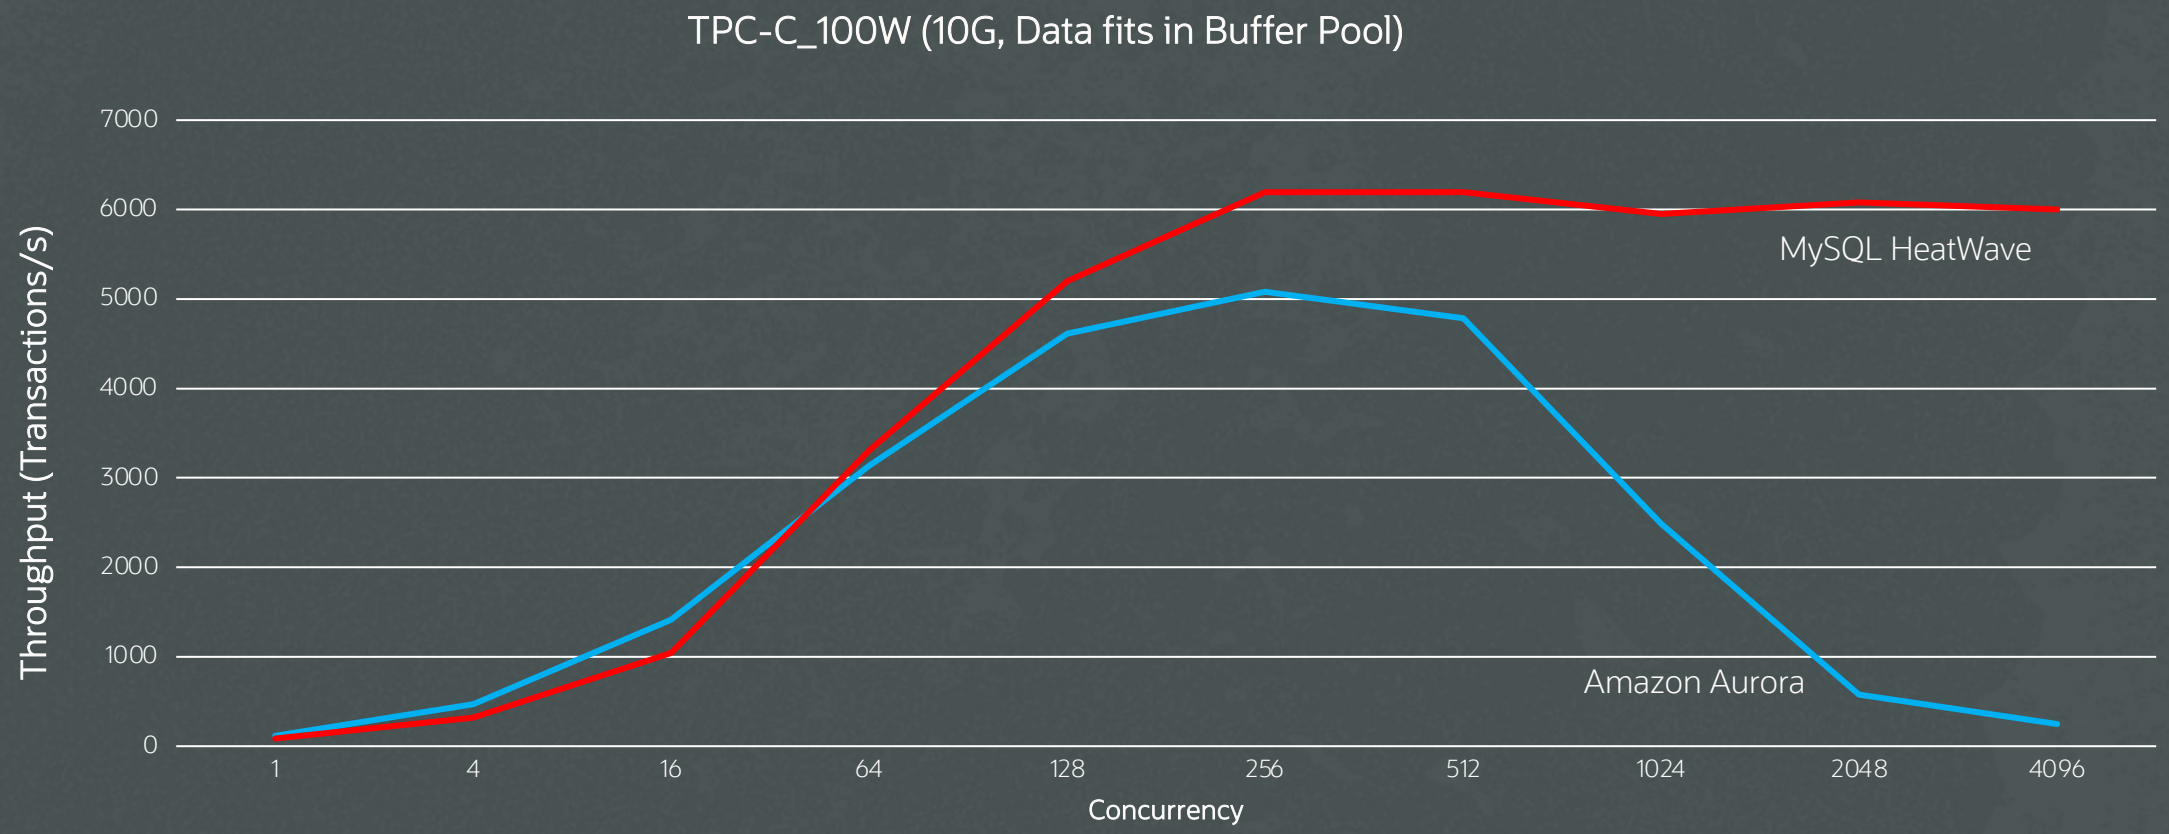

*Benchmark queries are derived from the TPC-C benchmarks, but results are not comparable to published TPC-C benchmark results since these do not comply with the TPC-C specifications.

Unbeatable price/performance for OLTP

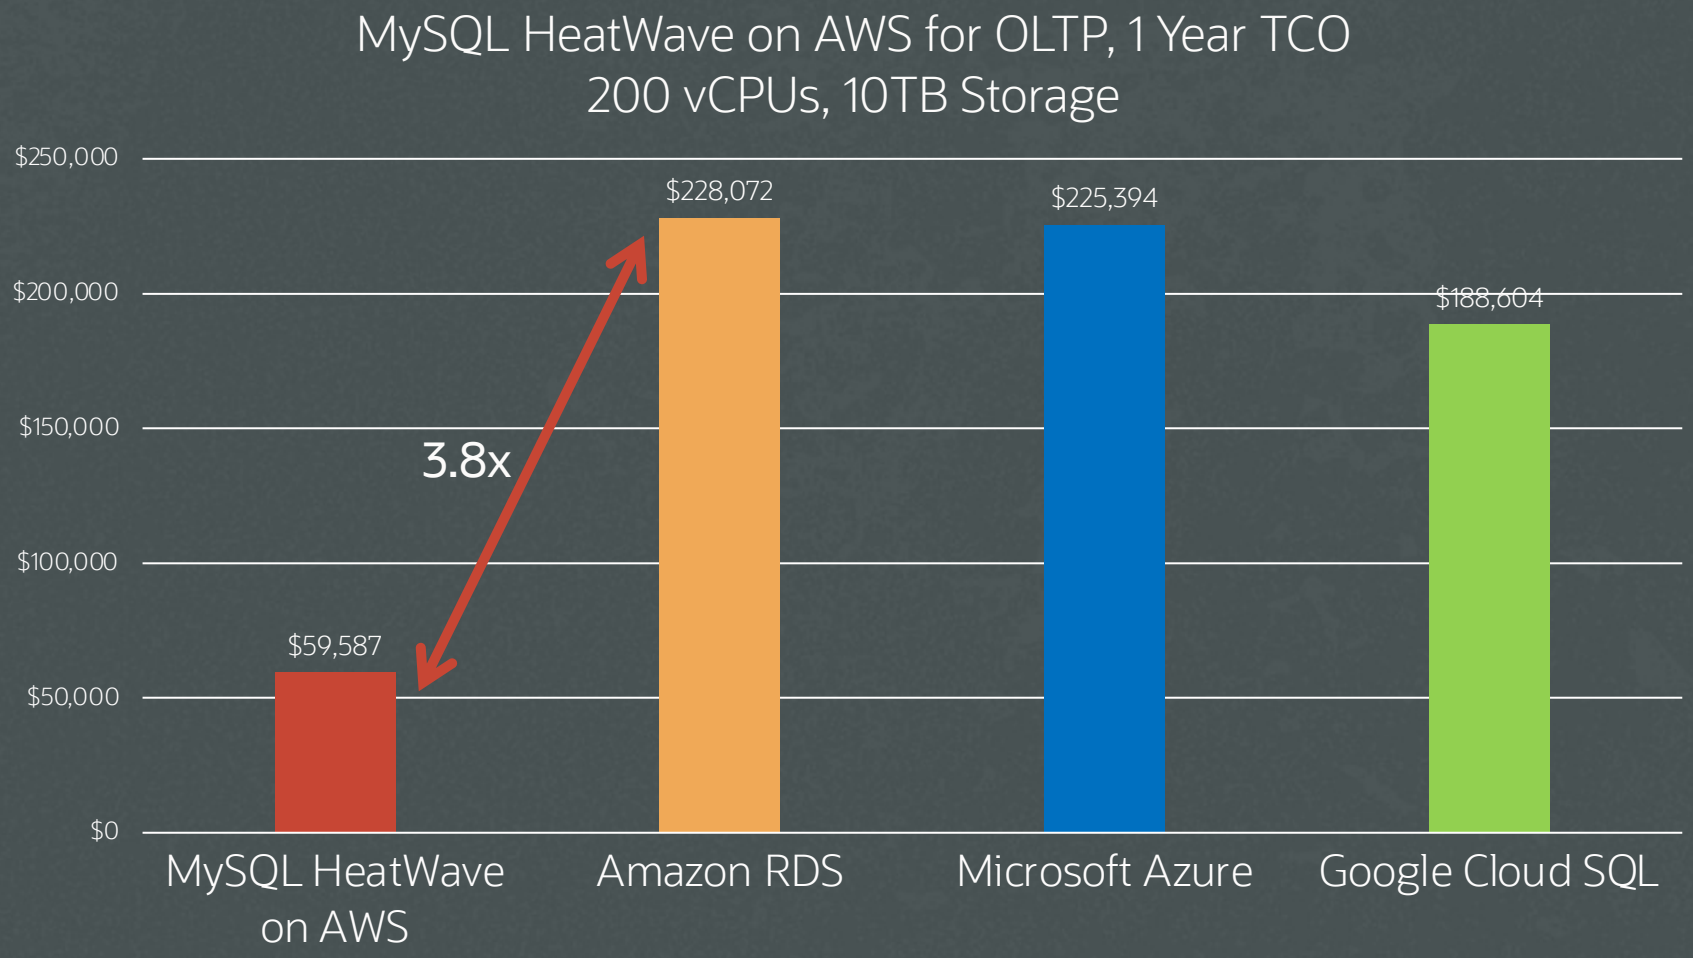

MySQL HeatWave: 1 ECPU, all regions have the same price.

Amazon RDS: Intel R5 16GB/Core, AWS US East.

Azure: Memory Optimized Intel 20GB/Core, MS Azure US-East.

Google: High Memory N1 Standard Intel 13GB/Core, GCP Northern Virginia.

Unbeatable price/performance for small OLTP AWS RDS/Aurora shape instances

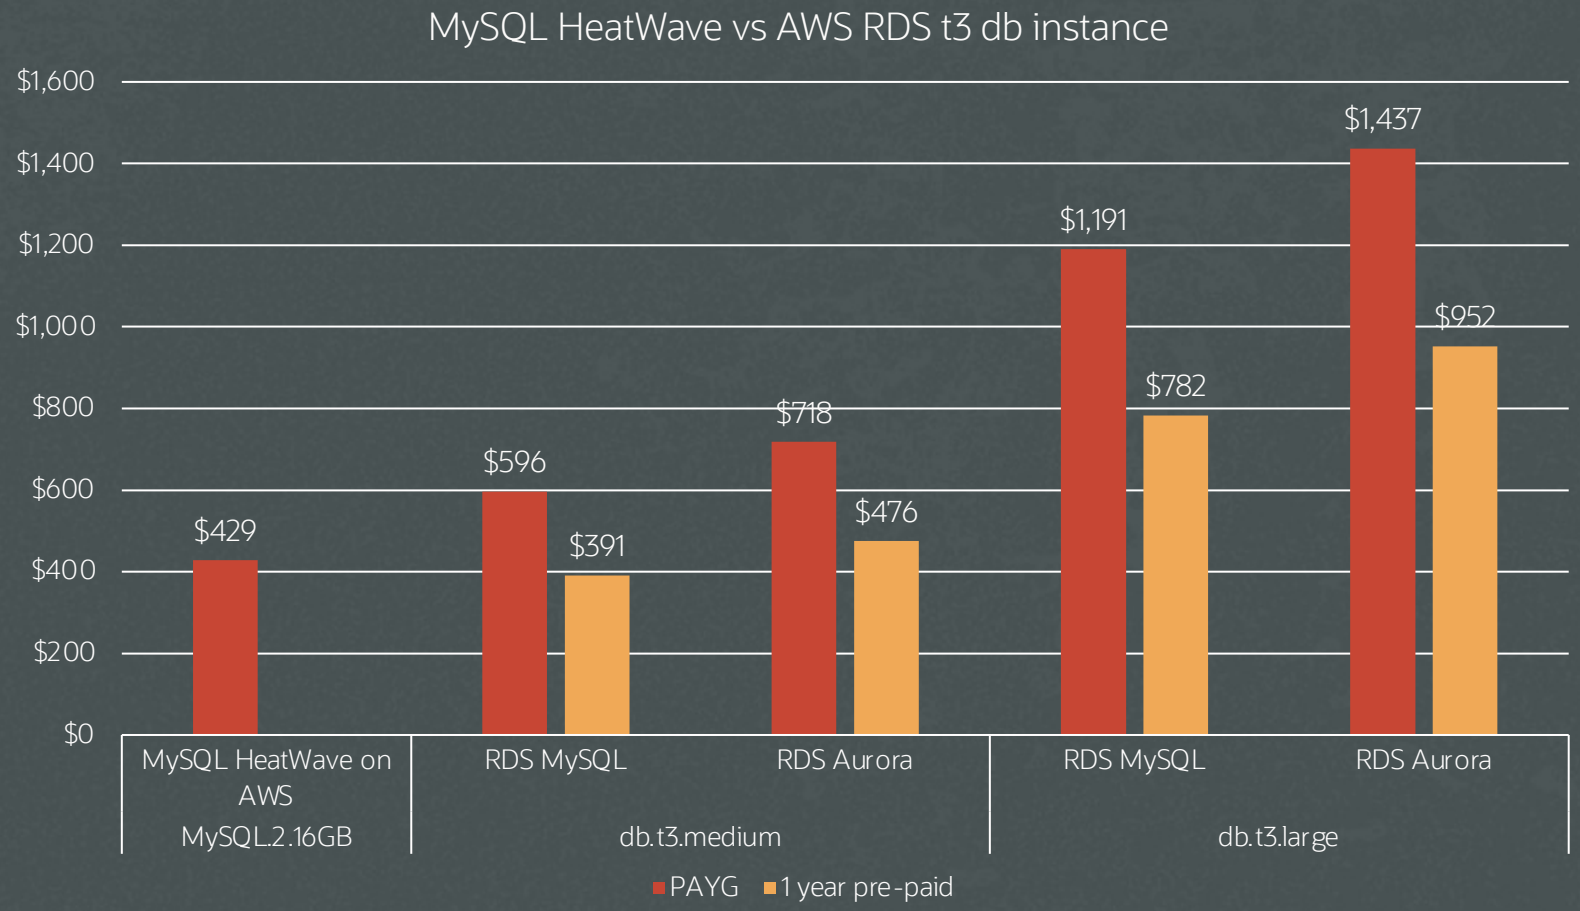

- 1.39x cheaper than RDS t3.medium
- 1.67x cheaper than Aurora t3.medium
- 2.78x cheaper than RDS t3.large
- 2.35x cheaper than Aurora t3.large

Note:
MySQL.2.16GB (2 vCPU, 16GB MEM)
RDS t3.medium (2 vCPU, 4GB MEM)
RDS t3.large (2 vCPU, 8GB MEM)

HeatWave AutoML: In-database machine learning

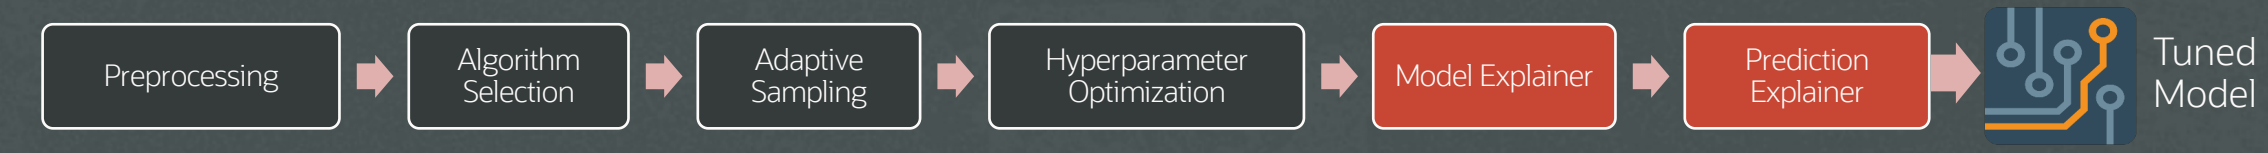

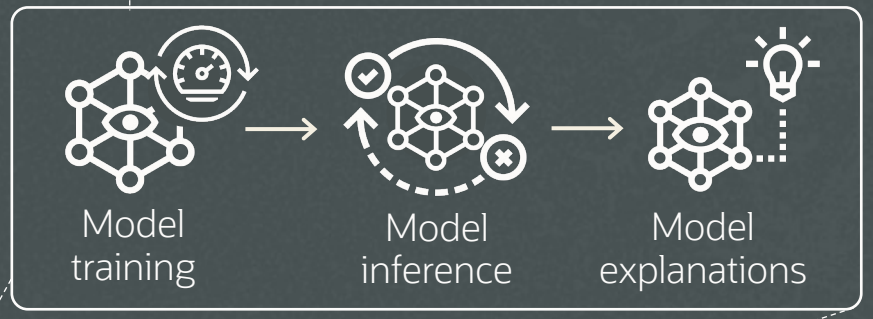

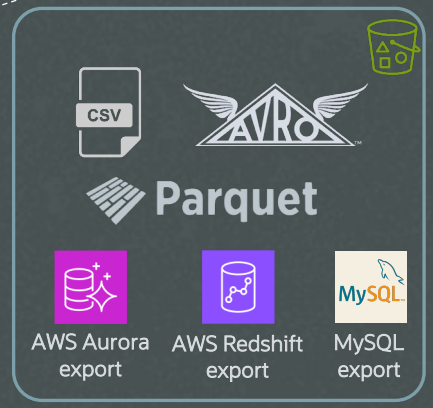

- Eliminates tedious and laborious steps
- Simple to use interface for beginner or advanced ML users
- Automatically selects algorithm and tunes it
- Explainable model behavior and predictions
- Fast training allows to quickly iterate to achieve desired outcome

Industries and use cases with HeatWave AutoML

Digital Marketing

Cost per acquisition

Targeted campaigns

Customer classification

E-Commerce

Videos for users

Lottery suggestions

Product upsell

Education

Predict student success

Monitor student behavior

HIPPA Compliance

Services

Erroneous ledge entries

Predict future losses

Predict price elasticity

FinTech

Loan default prediction

Identify loan extensions

Loan approval

Gaming

Player churn detection

Adjust game difficulty

Identify game hackers

Internet Of Things

Airport ticketing

Rain water level

Air pollution

Manufacturing

Reduce warranty claims

Defective part identification

Detect anomalies in supplies

Machine learning with HeatWave is fast, cost effective, accurate and sclable

25x
faster than Redshift

1%
of the cost of Redshift

MySQL HeatWave on AWS Console

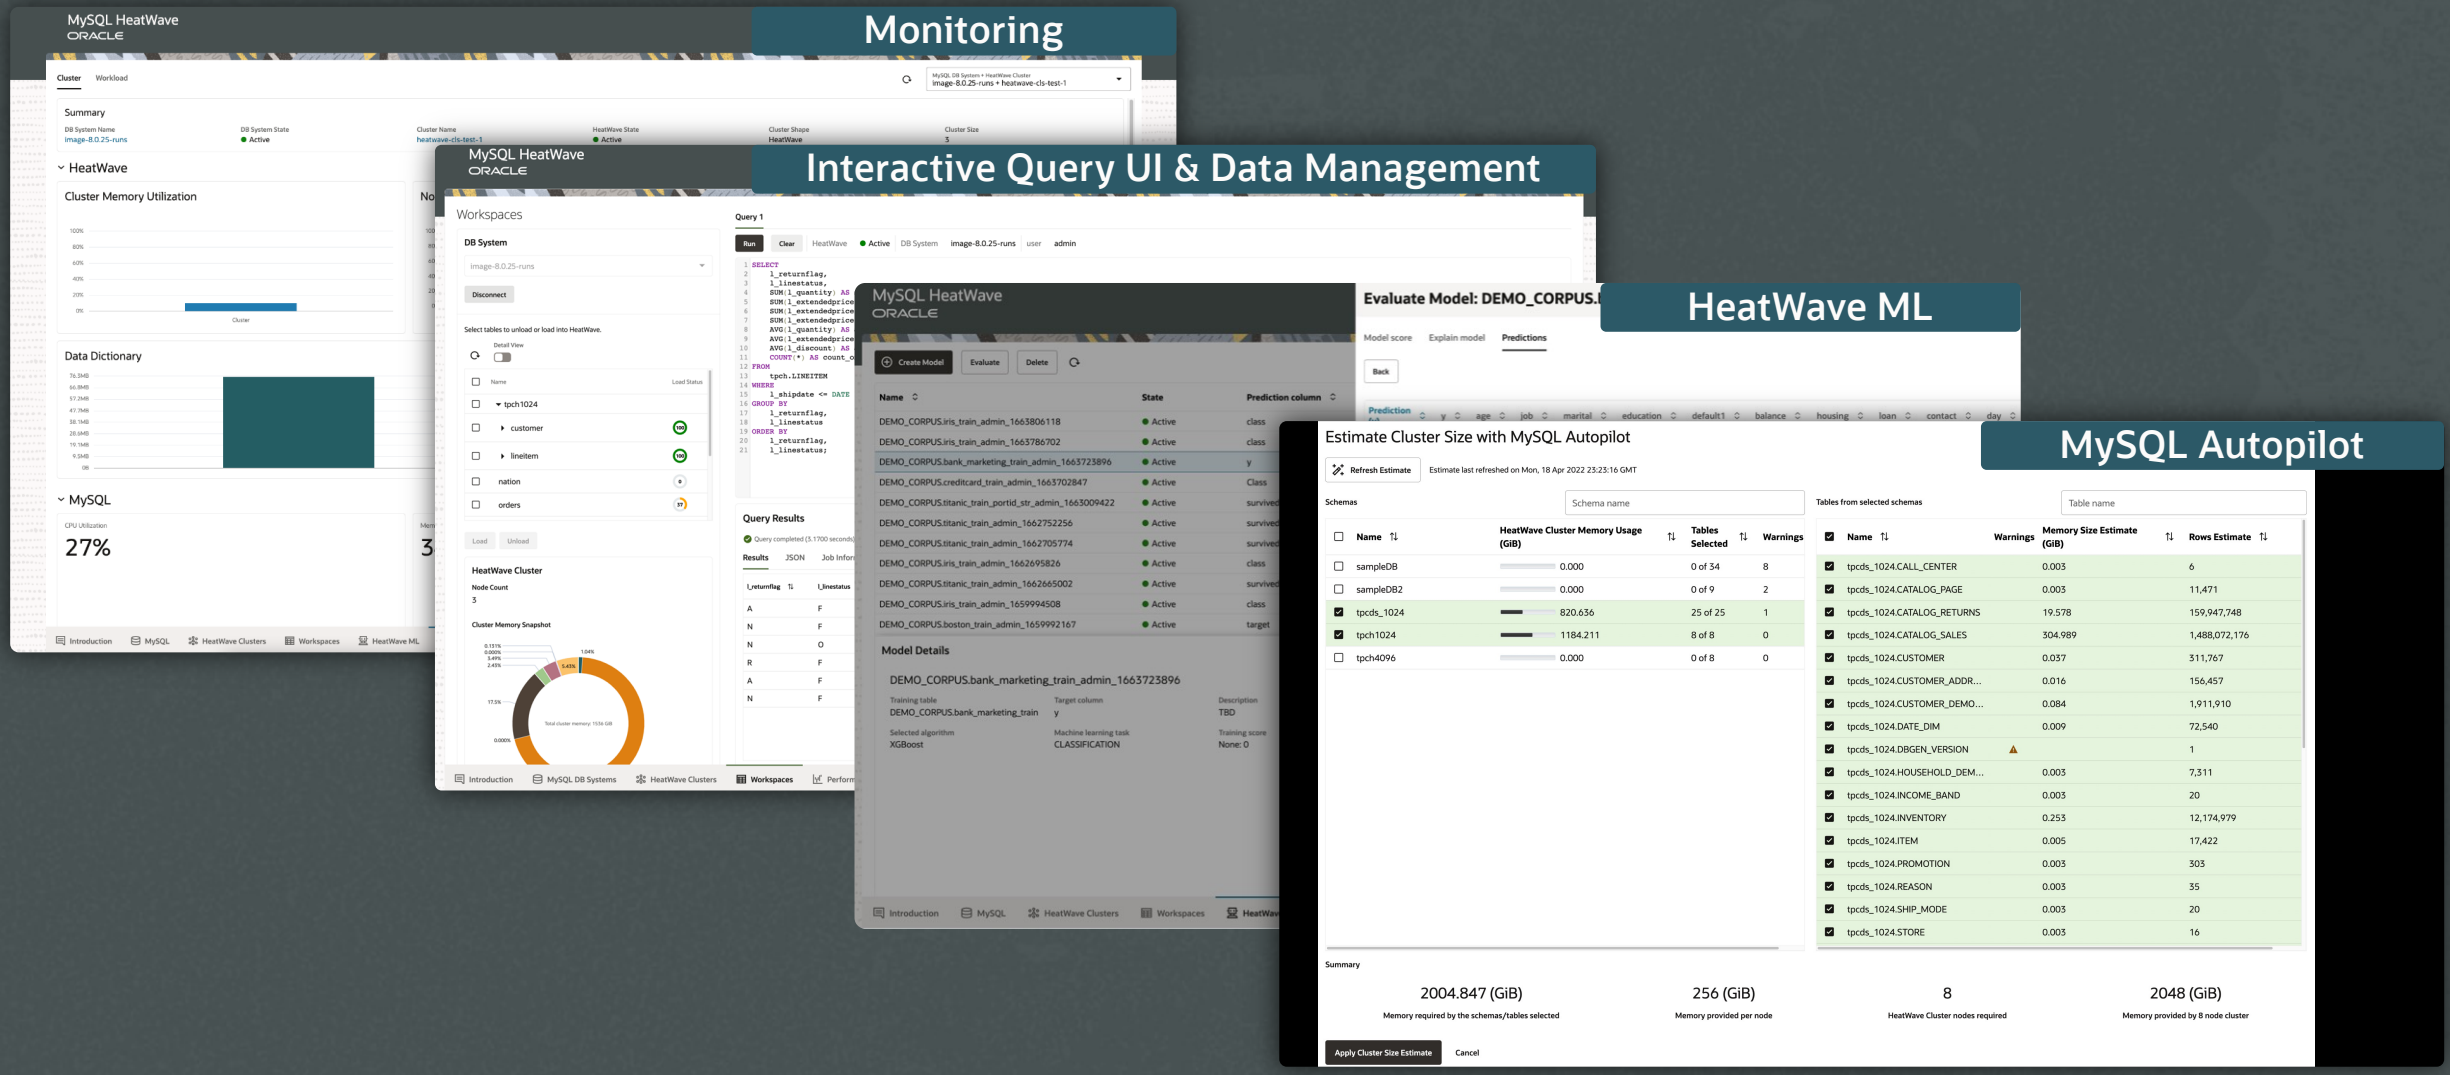

Quick start with a few clicks

Starter DB System

- Create MySQL HeatWave instance with predefined configurations and sample data pre-loaded to HeatWave

Sample Data

- Import sample schema such as airportdb and TPCH into the DB System
- Run sample analytics queries
- Evaluate MySQL HeatWave performance and features

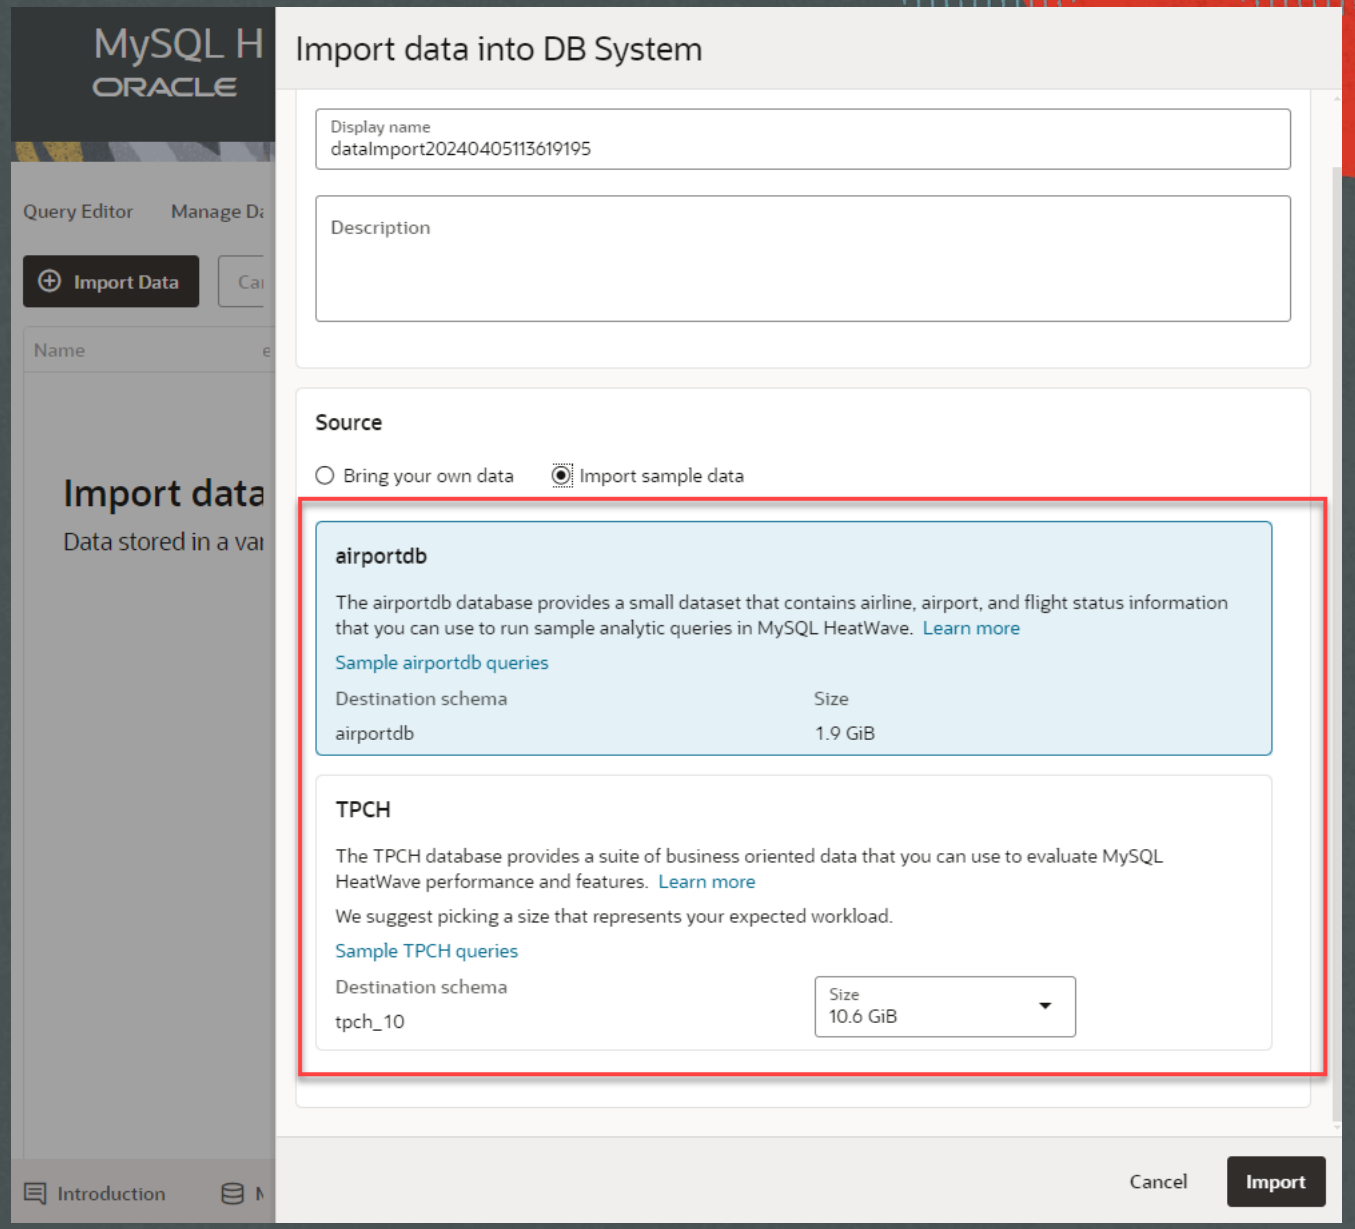

HeatWave AutoML console

Democratizing machine learning

- Visual interface
- Enables business analysts to build, train, run, and explain ML models. No SQL and no coding required
- Ability to explore “What-if” scenarios
- All processing is done inside the database
- No additional cost to users

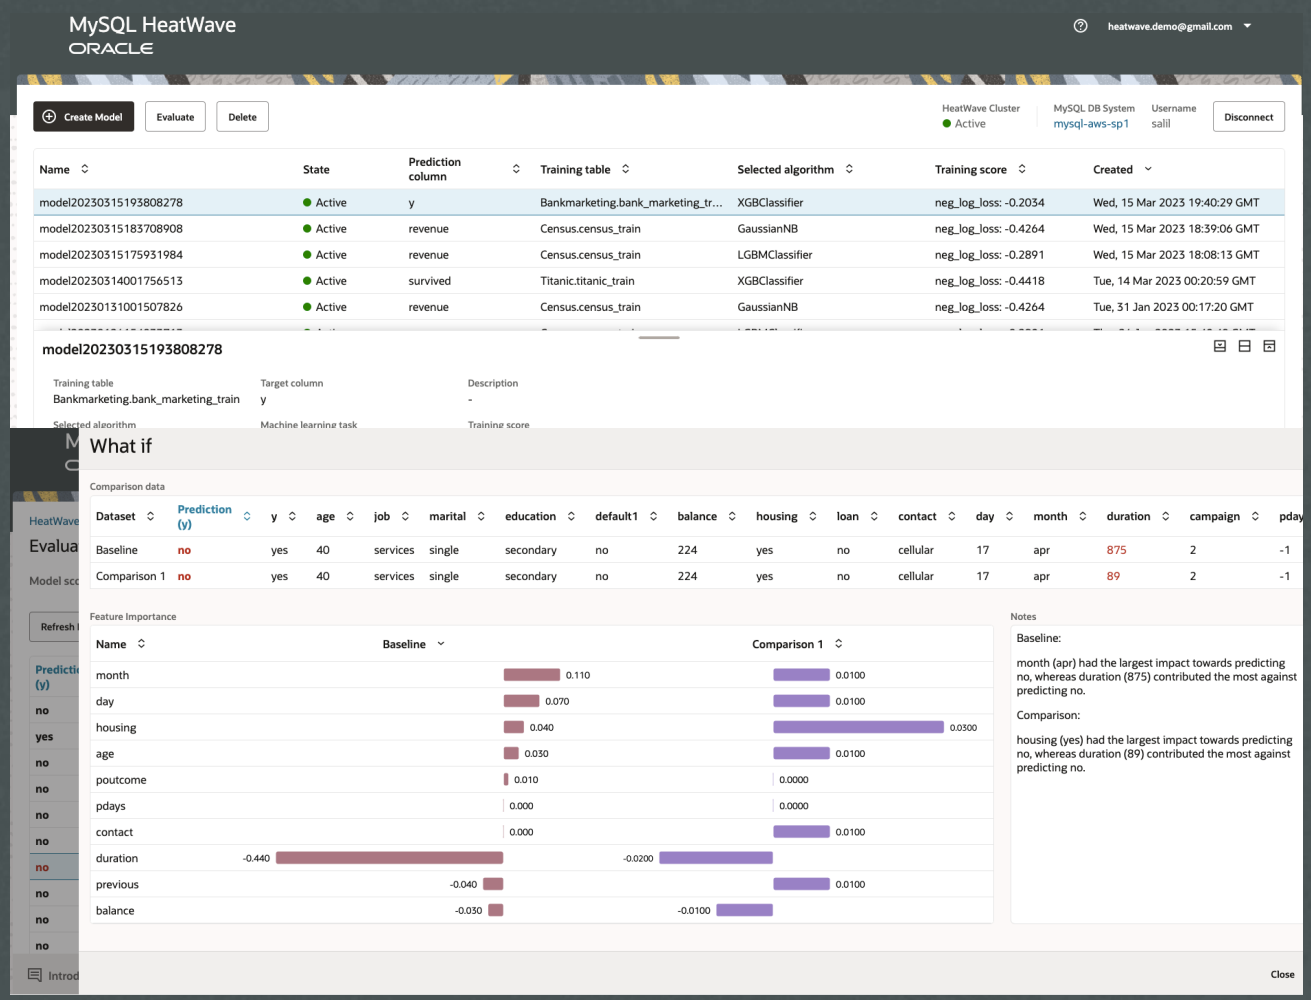

Aiwifi migrates from Amazon RDS to MySQL HeatWave on AWS

"With MySQL HeatWave's incredible performance and built-in machine learning, we knocked down previous barriers to growth. Aiwifi estimates that MySQL HeatWave replaced up to 5 external systems. Making HeatWave available on multiple cloud platforms is a very smart move by Oracle."

Eric Aguilar, CEO & CTO, Aiwifi

Challenge:

Aiwifi is a Mexican company developing Wi-Fi solutions that connect shoppers to websites through customized captive portals. Its value proposition is to gather valuable customer data by tracking user profiles and activity. Upon starting up in 2019, Aiwifi chose AWS as its platform and Amazon RDS as the backend database. However, as the business rapidly grew and generated heavy data loads, the lack of performance became a bottleneck for sustained growth, and database costs became a heavy challenge. In 2023, it migrated from Amazon RDS to MySQL HeatWave running natively inside AWS.

Results:

- ✓ Queries ran 13X faster and loading time on captive portals dropped by 50%, allowing Aiwifi to quickly onboard new customers without added costs
- ✓ Costs were reduced by 50%. The MySQL HeatWave high performance allowed using a smaller instance and high data egress fees were eliminated
- ✓ MySQL HeatWave efficiently handles complex queries on more than 40 million records to provide real-time analytics dashboards
- ✓ The need for query optimization was eliminated, allowing Aiwifi's developers to focus on building machine learning models with HeatWave AutoML
- ✓ ML is used to segment their user base and create more personalized marketing content as well as to predict offers that could be of interest to different customer segments

Products Used:

Oracle MySQL HeatWave

Read story

Software company utilizing ML/GenAI

Company: Provides Automation of ITSM & GRC in an Integrated Platform

Use Cases:

1. Pre-configured processes and workflows eliminating spreadsheets and manual work
2. Maximum visibility and data insights allows users to correlate, analyze, and remediate issues
3. Flexible Platform that can scale and simplify existing stack

Solution: Leveraging HeatWave to Automate IT & Security Management

50% Reduction in ML Activity

Reduction in Data Cleaning, Model Selection, Model Tuning and Training Time

15-25% Performance Improvement

Using Auto-Indexing, Auto-ML and Javascript
Move processing closer to Data to improve latency (Javascript)

Tech. Consolidation

Seamlessly extending relational model to support OLAP/ML/AI workload
Reduction in maintaining and deploying multiple Technology stack and training

Eat Easy
Food Delivery

Dubai-based online aggregator connecting thousands of its users to their favorite restaurants, making online ordering easier, reliable and convenient.

Use Case

- Predict food deliver time.
- Suggest food / restaurant based on past actions.
- Summarize menu for the selected restaurant.

Model Type

- Regression, Recommendation, Generative AI

Results

- Developed ML models in days that would have otherwise taken months
- Database developers were able to build the models without ML expertise
- Simplified infrastructure with no complex ETL to manage and one platform providing OLTP, Analytics, ML and GenAI
- Consistent interface across various ML model types simplified learning for Eat Easy development team

MySQL HeatWave on AWS

Combine FIVE AWS services into ONE

- MySQL HeatWave runs natively on AWS, optimized for AWS infrastructure
- Data doesn’t leave AWS – no data egress costs, and avoids compliance approvals
- Lowest latency access to MySQL HeatWave
- Tight integration with the AWS ecosystem – S3, CloudWatch, PrivateLink
- Easier migration from other databases (e.g., Amazon Aurora, Redshift, Snowflake)