

State of the Dolphin

Wim Coekaerts

Executive Vice President, Software Development, Oracle

May 1, 2024



Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Agenda

Data Platform Complexity

Technology Update

- Generative AI and Vector Store
- MySQL HeatWave
- Server Improvements
- Accelerated Development
- Observability & Management




Data Platform Complexity

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

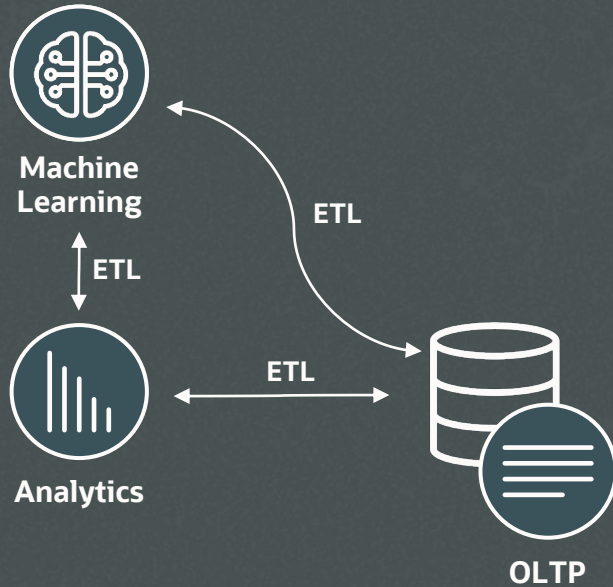


Data Platform Complexity

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

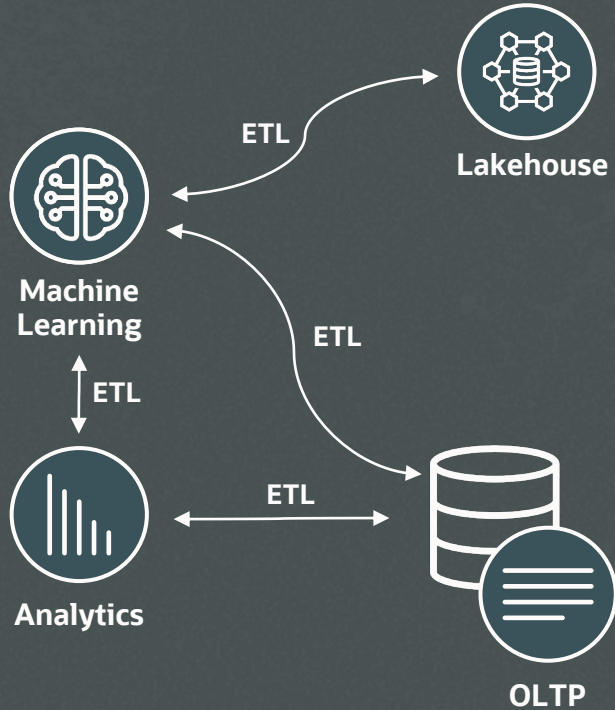


Data Platform Complexity




- **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



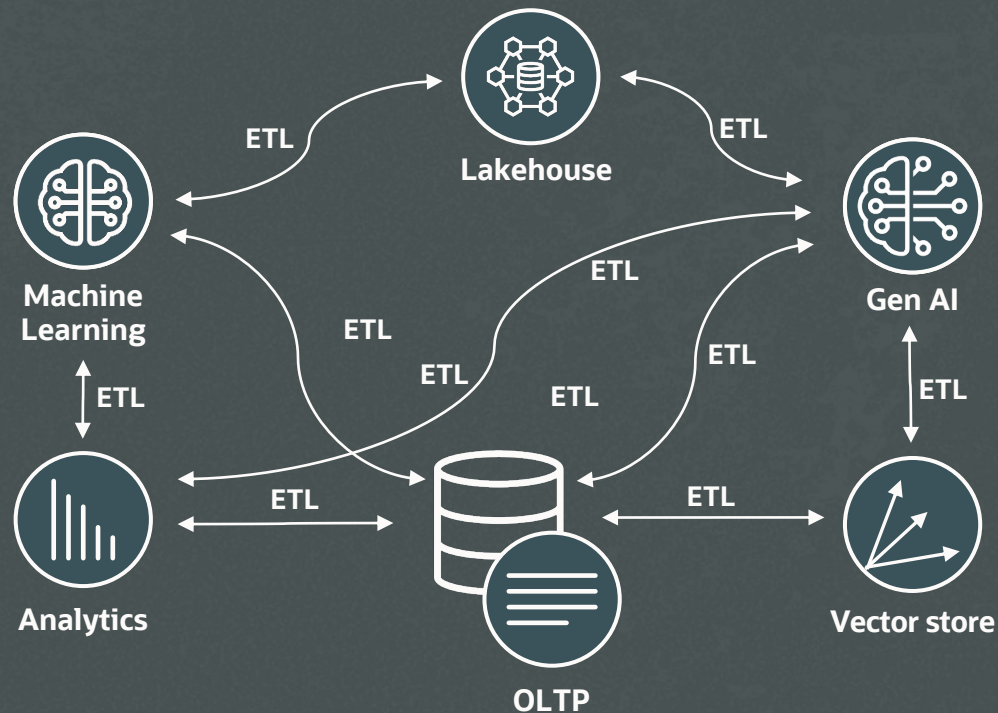
- **You start** with an OLTP database application
- **Analytics** will give LOB managers valuable insights
- **Machine learning** predictions will improve the customer experience
- **Lakehouse** will deliver insights into structured and semi-structured data

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 structured and semi-structured 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 structured and semi-structured 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 exposure
- Requires specialized skills

Data Platform Complexity at Oracle (OCI)



- ✓ Gen AI
- ✓ Vector store
- ✓ Lakehouse
- ✓ OLAP
- ✓ OLTP

One Database Cloud Service

- Democratizes OLAP, ML, AI
- No complex ETL
- Real-time data
- Secure
- Easy to manage

MySQL HeatWave

ONE DATABASE FOR OLTP, OLAP, LAKEHOUSE AND ML

Social, eCommerce, gaming, healthcare, fintech apps. Analytics and ML tools



ORACLE
Analytics Cloud



Queries

Results

MySQL HeatWave

Real-time analytics,
machine learning,
lakehouse, and OLTP in
one cloud database
service



OLTP



Analytics



AutoML



Autopilot



Lakehouse



Technology Update



Generative AI in HeatWave Enables New Use Cases




Retrieval Augmented Generation (RAG)

- Generate insights from enterprise documents
- Generate blogs from pdf instruction manuals
- Summarize logs



Content generation & summarization

- Search on public and private enterprise data
- Search on unstructured data in vector store



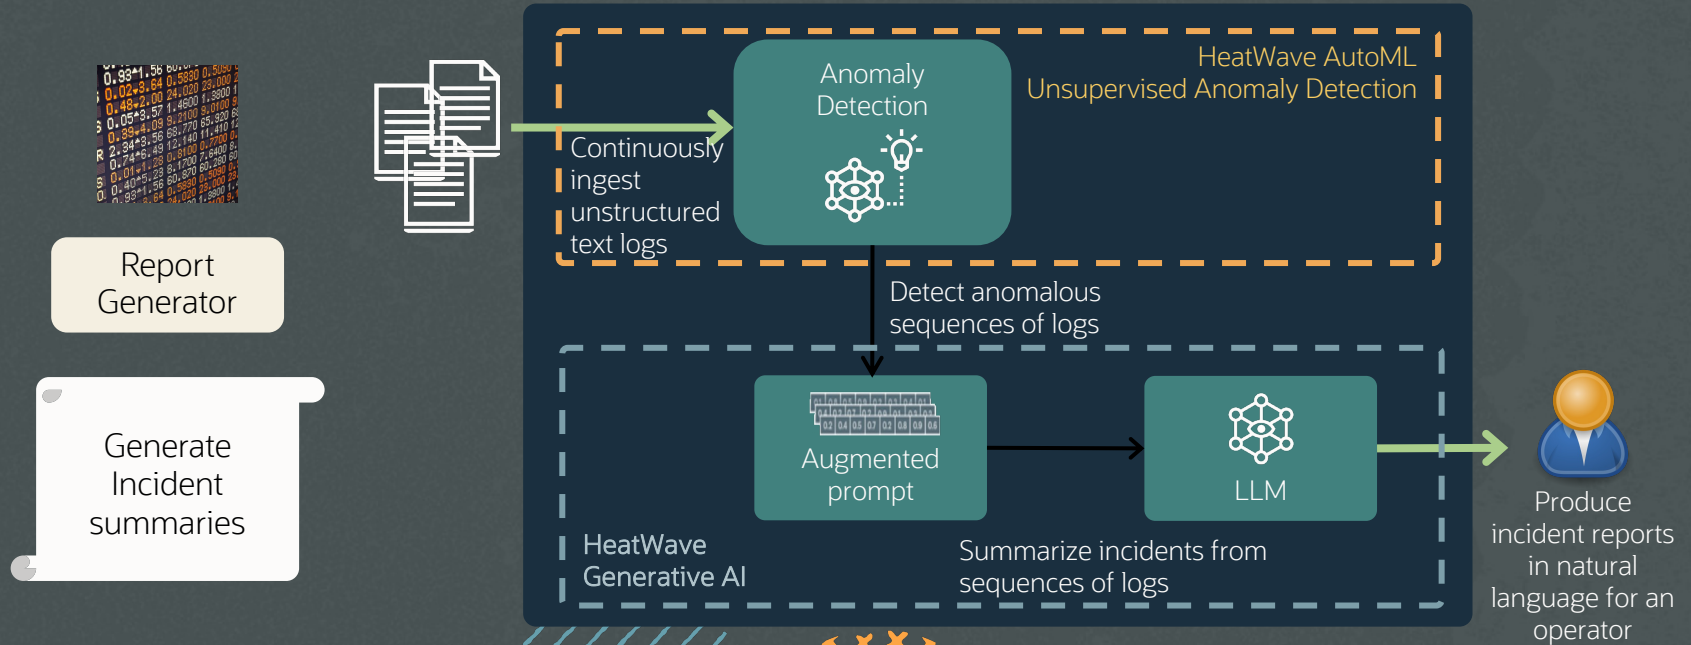
Natural language interaction

- Natural language interaction with unstructured data
- Content retrieval and response in natural language

Usecase 1: Summarization

CONTENT GENERATION/SUMMARY

Report Generation: HW AutoML + Summarization of anomalous logs



Usecase 1: Summarization

SYNERGY BETWEEN AUTOML AND GENERATIVE AI

Report Generation: Summarization of anomalous logs

```
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777883620", "msg": "[8145-140643603412800] [fit_transform:40] [ 1.939] [fit_transform:Remove non-categorical columns with >20% missing values. Dropping 0 columns]"
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777883805", "msg": "[8145-140643603412800] [fit_transform:147] [ 1.939] [fit_transform:Remove constant columns. Dropping 0 columns]"
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777964868", "msg": "[8145-140643603412800] [fit_transform:227] [ 173.397] [fit_transform:Alter feature engineering and transformations. Updated shape : [157776, 72967]}"
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777964868", "msg": "[8145-140643603412800] [fit_transform:231] [ 173.397] [fit_transform:col_types : {'text': 'text'}]"
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777964879", "msg": "[8145-140643603412800] [fit_transform:247] [ 173.397] [fit_transform:class counts : 0 157494\\n1 282\\n\\nName: label, dtype: int64]"
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975619", "msg": "[8145-140641670100736] [run:272] [ 249.292] [249.29190063476562, 249.29190063476562] (current,max) GB for process 8145 is NOT within the threshold 248.94811630249023 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975619", "msg": "[8145-140641670100736] [run:279] [ 249.292] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975725", "msg": "[8145-140641670100736] [run:272] [ 250.038] [250.03799438476562, 250.03799438476562] (current,max) GB for process 8145 is NOT within the threshold 248.19862747192383 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975725", "msg": "[8145-140641670100736] [run:279] [ 250.038] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975830", "msg": "[8145-140641670100736] [run:272] [ 250.782] [250.7821350076562, 250.7821350076562] (current,max) GB for process 8145 is NOT within the threshold 247.45304489135742 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975830", "msg": "[8145-140641670100736] [run:279] [ 250.782] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975936", "msg": "[8145-140641670100736] [run:272] [ 251.528] [251.5282275976562, 251.5282275976562] (current,max) GB for process 8145 is NOT within the threshold 246.70550918579102 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777975936", "msg": "[8145-140641670100736] [run:279] [ 251.528] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976042", "msg": "[8145-140641670100736] [run:272] [ 252.274] [252.2743250976562, 252.2743250976562] (current,max) GB for process 8145 is NOT within the threshold 245.95797729492188 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976042", "msg": "[8145-140641670100736] [run:279] [ 252.274] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976148", "msg": "[8145-140641670100736] [run:272] [ 253.02] [253.0204162976562, 253.0204162976562] (current,max) GB for process 8145 is NOT within the threshold 245.21044158935547 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976148", "msg": "[8145-140641670100736] [run:279] [ 253.02] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976253", "msg": "[8145-140641670100736] [run:272] [ 253.767] [253.7668100976562, 253.7668100976562] (current,max) GB for process 8145 is NOT within the threshold 244.46290588378906 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976253", "msg": "[8145-140641670100736] [run:279] [ 253.767] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976359", "msg": "[8145-140641670100736] [run:272] [ 254.513] [254.51260375976562, 254.51260375976562] (current,max) GB for process 8145 is NOT within the threshold 243.71537017822266 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976359", "msg": "[8145-140641670100736] [run:279] [ 254.513] Sending SIGTERM to the main thread"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976465", "msg": "[8145-140641670100736] [run:272] [ 255.261] [255.26065063476562, 255.26065063476562] (current,max) GB for process 8145 is NOT within the threshold 242.96588134765625 GB memory usage"}
[analyticsCluster]: {"component": "ml_driver", "level": "INFO", "ts": "1688777976465", "msg": "[8145-140641670100736] [run:279] [ 255.261] Sending SIGTERM to the main thread"}
```

Ask the question (“What is the main problem in the following collection of logs. Provide a two-sentence summary.”)

“The main problem in this collection of logs is that the memory usage of a particular process (with ID 8145) is consistently exceeding the defined memory threshold. This issue is repeatedly highlighted in the logs, and SIGTERM signals are sent to the main thread to terminate the process when it exceeds the memory threshold. Would you like me to help you with anything else regarding this collection of logs?”



HeatWave Offers Choice for Running LLMs

In-HeatWave LLMs

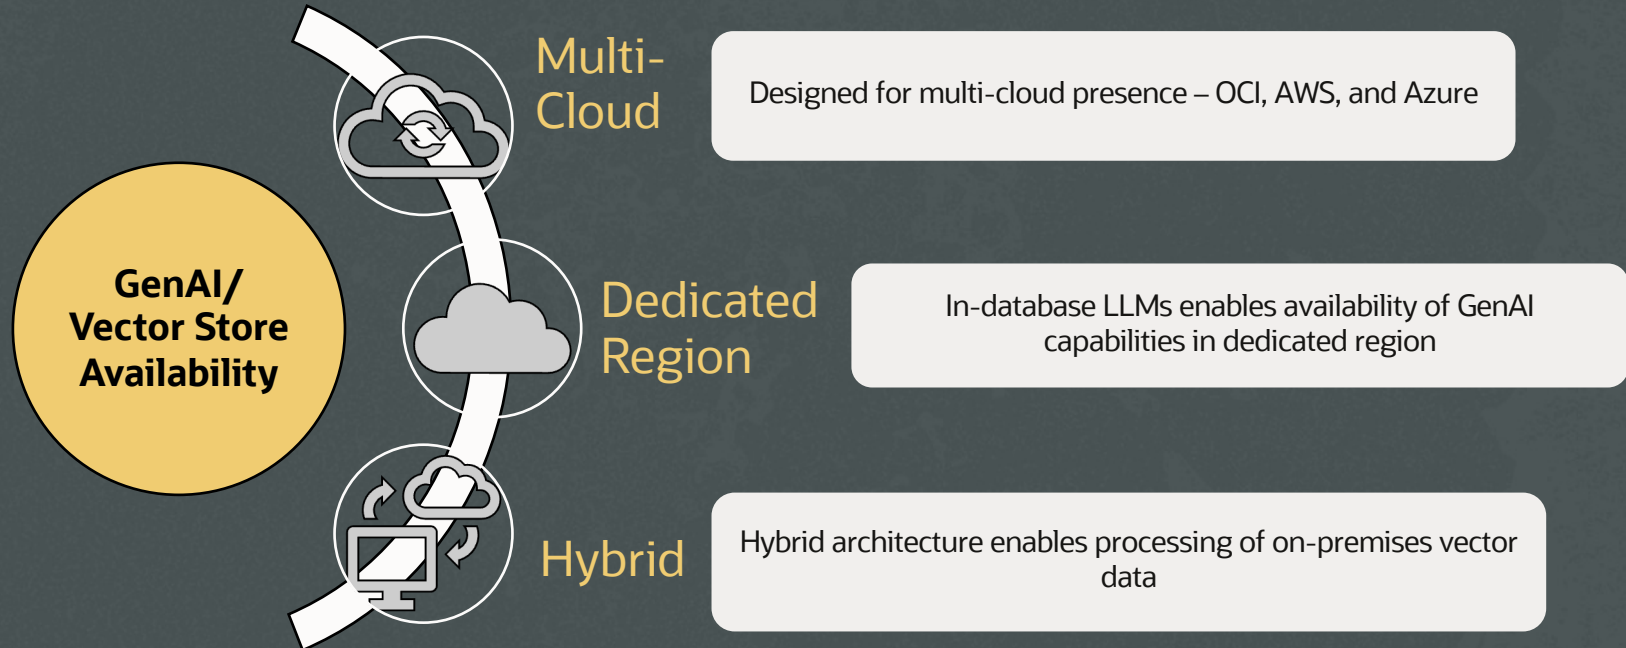
- Native execution within the HeatWave database
- Run smaller LLMs like Llama2-7B and Mistral-7B
- Secure, lower cost, guaranteed availability

OCI Generative AI service LLMs

- Support larger models like Cohere-command and Llama2-70B and run on GPUs
- Higher quality, better performance

Available in All Regions

VECTOR STORE AND ABILITY TO RUN LLMS INSIDE HEATWAVE PROVIDES FLEXIBILITY OF DEPLOYMENT



Vector Data Type Support

Standard SQL interface to create tables with vector columns

Vector data storage

- HeatWave: In-memory columnar format
- InnoDB: BLOB


```
mysql> CREATE TABLE wikipedia (  
    id INT,  
    title VARCHAR(1024)  
    page_data TEXT  
    page_list TEXT,  
    page_url TEXT,  
    page_embedding VECTOR(1024)  
ENGINE_ATTRIBUTE={'model':"cohere"})  
ENGINE=lakehouse, SECONDARY_ENGINE=rapid;
```

Example distance functions

- L1/MANHATAN
- L2/EUCLIDIAN
- L1^2/MANHATAN_SQUARED
- L2^2/EUCLIDIAN_SQUARED
- COSINE
- DOT
- HAMMING

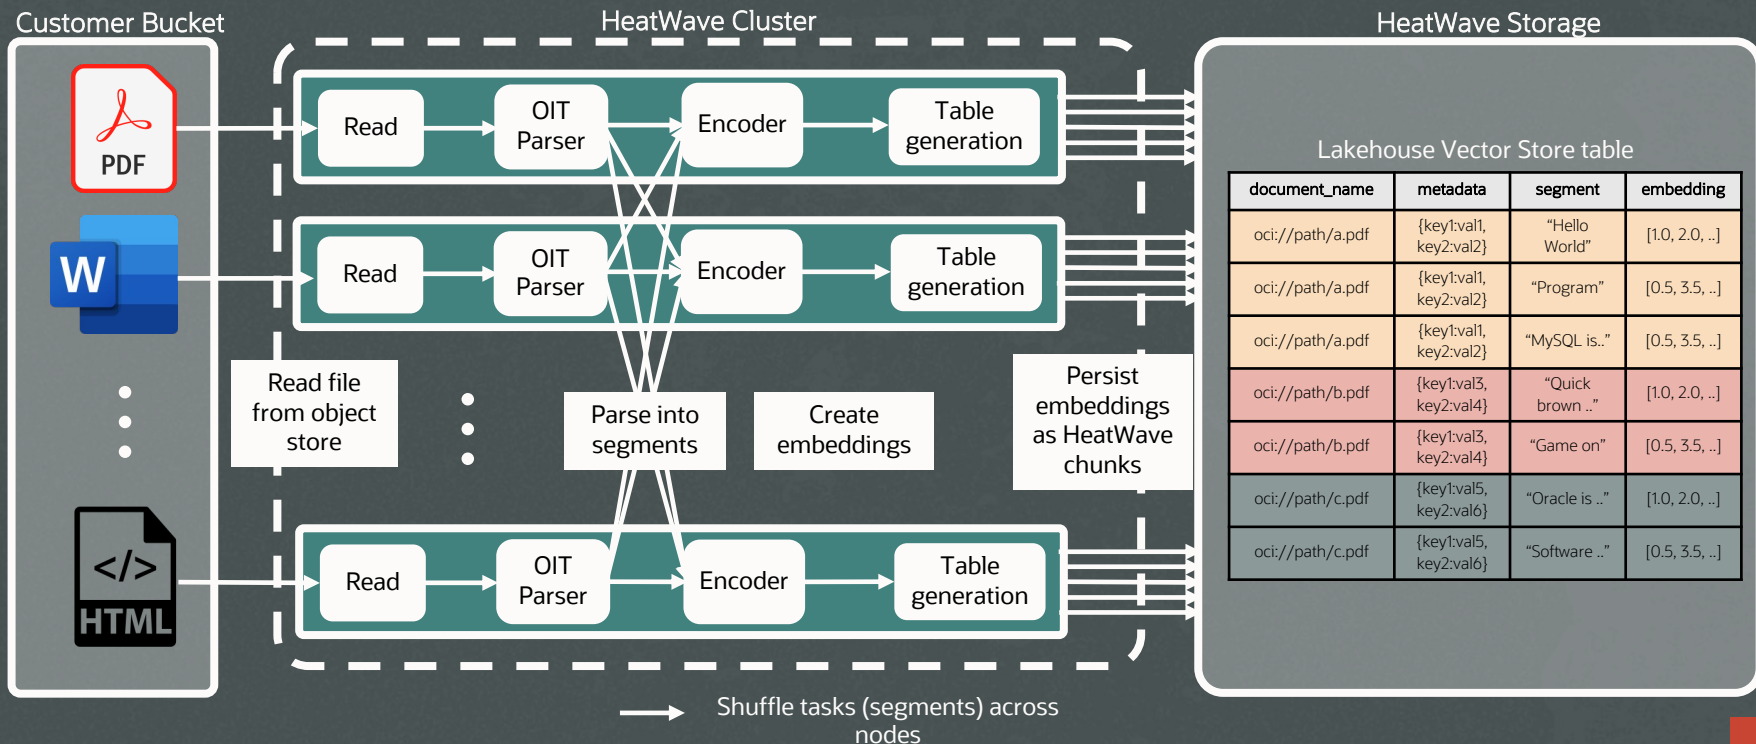


All phases of Creating a Vector Store Done Inside HeatWave



Vector Store Creation with HeatWave is Parallelized

BUILT IN EMBEDDING GENERATION DONE ACROSS NODES



DEMO: Gen AI

Sandeep Agrawal



MySQL HeatWave



OLTP with MySQL HeatWave

High Availability

Made up of three MySQL instances: a primary instance and two secondary instances

Provisioned across different availability or fault domains

The primary instance functions as a read/write endpoint

Automatic or Manual Promotion of a Secondary Instance

Automatic connection failover with zero data loss

Scalability

Read replicas for read-intensive workloads

Add and remove Read Replicas for horizontal read elasticity

Supports up to 18 read replicas

Built-in Load Balancer for the read-only endpoint

5-Tuple Hash Load Balancing Policy

Enhanced bulk ingest for faster imports from object store

Data protection and durability

Automatic backups

- Retention period between 1 and 35 days

Manual backups

- Retention for up to 365 days

Operator backup

- Created by CloudOps to assist in investigating potential issues

Point-In-Time Recovery

- RPO of approximately five minutes for an active database system
- RPO of 24 hours with daily backup

Reliability

Automated upgrades and patching

- No need to schedule patches, track patch contents or request one-off patches

Always up-to-date security fixes

- Frequent patches to meet compliance requirements

Full-stack patching

- Includes database and all cloud infrastructure



OLTP with MySQL HeatWave

SINGLE CLICK HIGH AVAILABILITY

- SLA 99.99%
- Automatic failover
- Zero data loss during failure event
- Disaster recovery with cross-region backups

Create DB System

Name: mysql-ocw-ha

Description: Optional

Standalone: Single-instance DB System

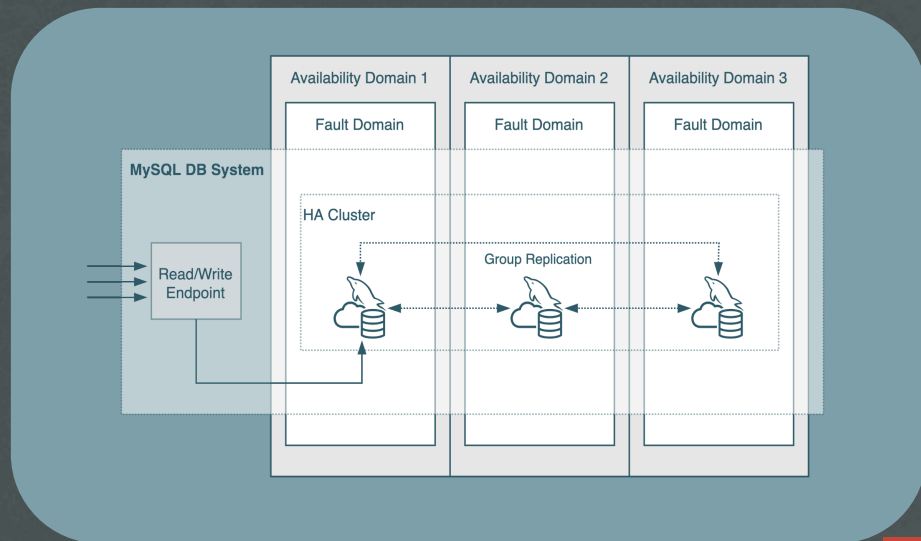
High Availability: Run a DB system with 3 MySQL instances providing automatic failover and zero data loss

HeatWave: DB System that allows you to enable HeatWave for accelerated query processing, suitable for running both OLTP and OLAP workloads

Create Administrator credentials

Username: admin

Buttons: Create, Cancel



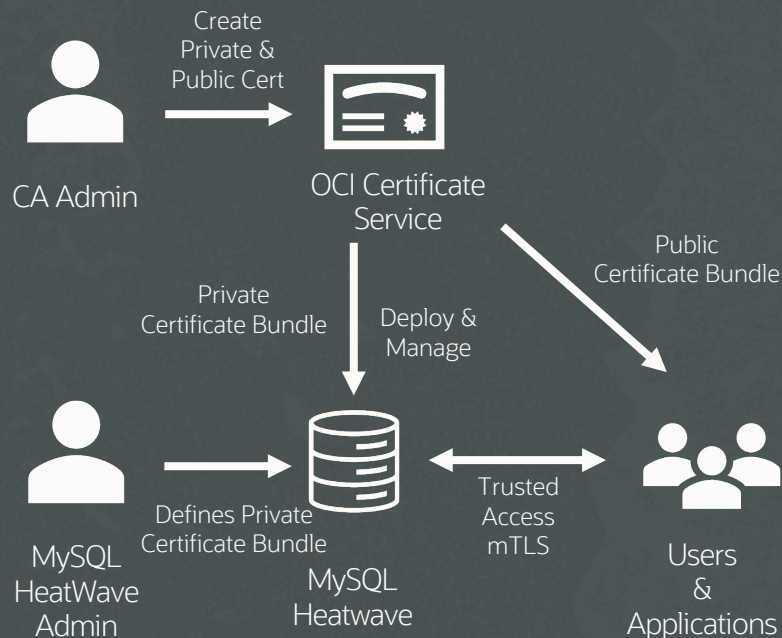
OLTP with MySQL HeatWave

BRING YOUR OWN CERTIFICATE (BYOC)

MySQL HeatWave enables users to point to the OCI Certificate Services so users can create, provision, manage, deploy, and automatically rotate PKI certs.

OCI Certificates Service

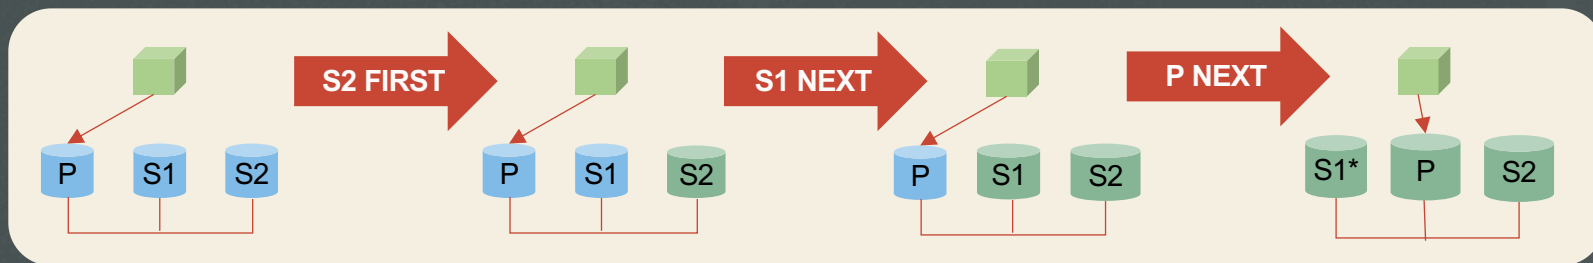
- Generates and stores private keys in a **FIPS Level 3 HSM**
- Removes the error prone, manual process of purchasing, uploading, and renewing certs
- Provides a Private Certificate Authority option for customers managing their own certificate chains
- Provides automatic cert deployment and renewal



OLTP with MySQL HeatWave

HA IMPROVEMENTS


- Improvements in downtime:
 - Planned Upgrade: **~60s** → **~15s**
 - Planned Switchover: **~60s** → **~7s**
- Automatically replace uncoverable instances
- Online changes to shape and configuration of a HA cluster



OLTP with MySQL HeatWave

READ REPLICA ENHANCEMENTS

Description	Summary
Change Shape	Customer can specify the shape for a specific read replica and can make it handle a specific read traffic profile that is more demanding resource-wise.
Change Configuration	Customer can specify the configuration for a specific read replica and can adjust it to handle a specific read traffic profile.
Change Version	Customer can specify the version for a specific read replica and try out new features with minimal disruption. Or they may want to pin a given replica to a specific version.



OLTP with MySQL HeatWave

PITR IMPROVEMENTS

1 5x faster, smaller vulnerability window

Binlog Size (GB)	Earlier	Now
4	52 sec	36 sec
16	10 min, 2 sec	2 min, 42 sec
44	15 min, 10 sec	2 min, 59 sec

2 Multi-Threaded PITR: 2x better performance

VM	Applying 500GB binlog	
	Earlier	Now
4 core VM	22 hours	10 hours
8 core VM	14 hours	7 hours

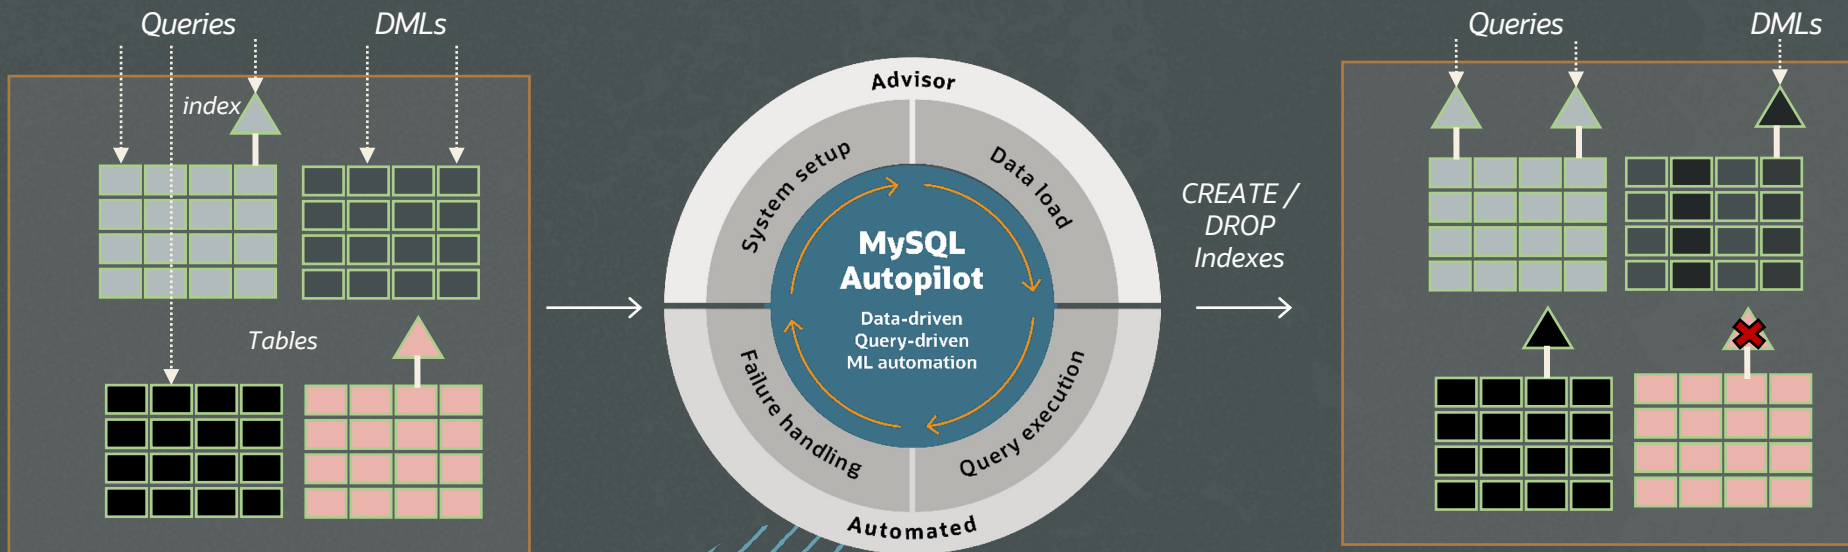
3 Improved disable PITR performance

- 34min to 3min
- Does not block other DB System operations

OLTP with MySQL HeatWave

MYSQL AUTOPILOT INDEXING (LA)

Recommends secondary indexes for OLTP workloads



New Operators Support in HeatWave


FACILITATES MIGRATION OF MORE WORKLOADS

Operator	Snowflake	Redshift	BigQuery	Databricks	PostgreSQL	MySQL HeatWave
CUBE	✓	✓	X	✓	✓	✓
HLL_COUNT	✓	✓	✓	✓	✓	✓
Grouping Sets	✓	✓	X	✓	✓	✓
Qualify	✓	✓	✓	✓	X	✓
Table Sample	✓	X	✓	✓	✓	✓



Best Performance in Industry even with Concurrency

TPC-DS* 100TB




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



Best Load Performance in the Industry

TPC-DS* 100TB



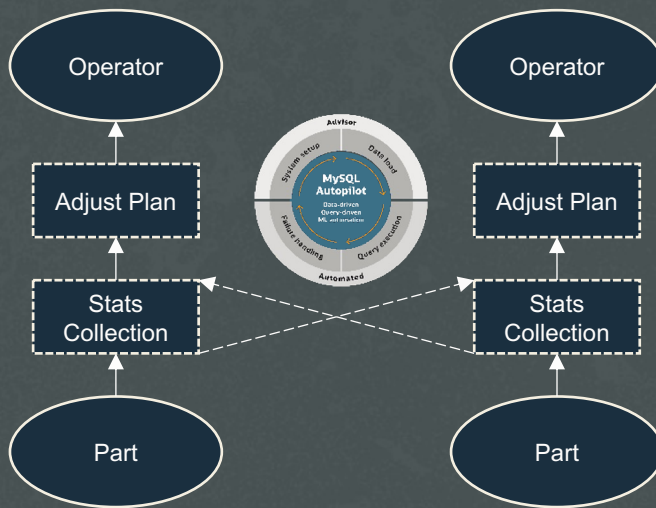
*Benchmark queries are derived from the TPC-DS benchmarks, but results are not comparable to published TPC-DS benchmark results since these do not comply with the TPC-DS specifications.



Adaptive Query Execution

- Dynamically adjusts data structures and system resources during query execution
- Independently optimizes query execution for each node based on actual data distribution at run time

Workload	Data size	Improvement in <u>first</u> run
TPCDS	2TB	21%
TPCDS	16TB	25%
TPCDS	100TB	10%

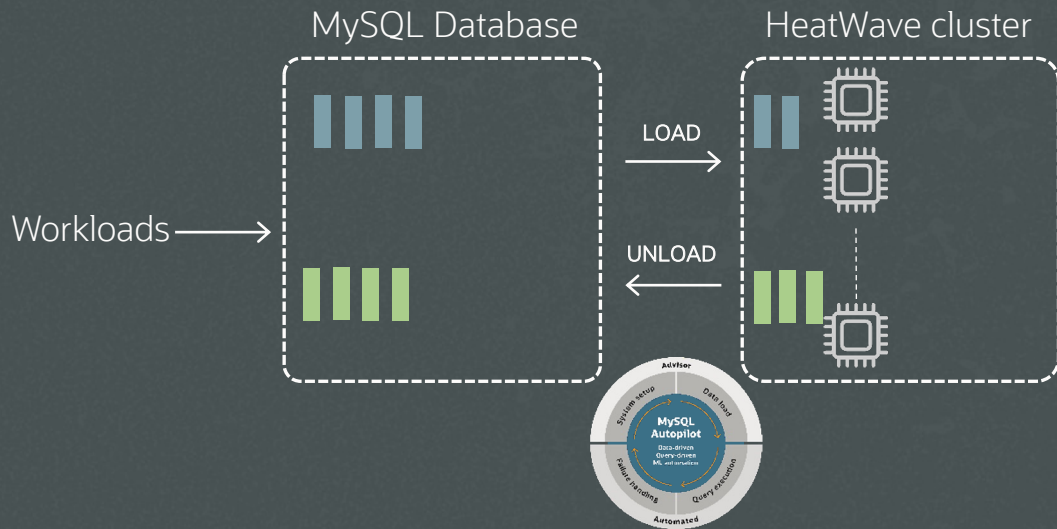


Improves ad-hoc query performance and skew handling



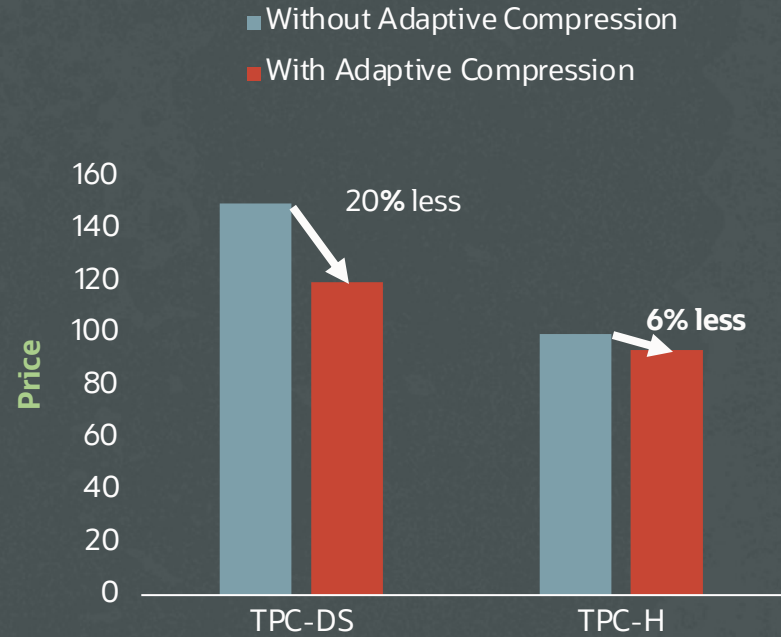
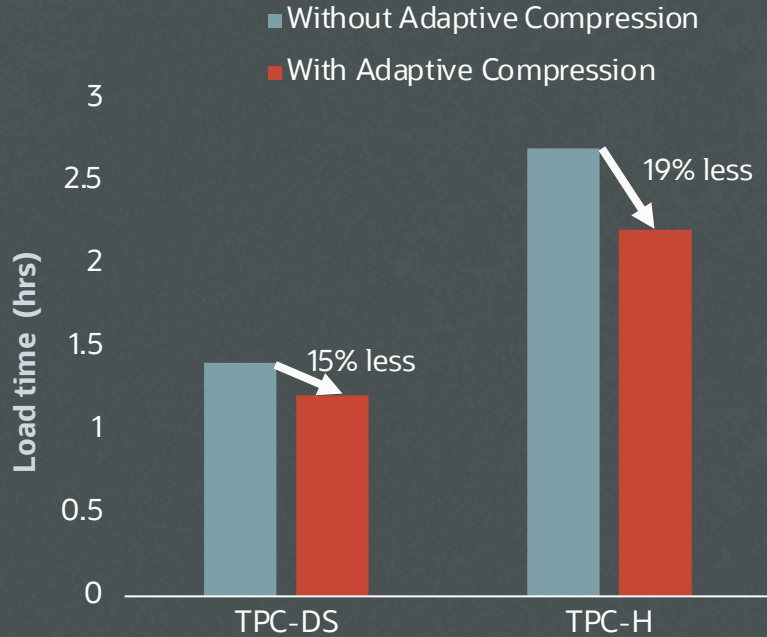
Auto Load and Unload

- Automatically loads tables or columns into HeatWave to optimize performance of user workload
- Automatically unloads tables less frequently used than other tables to optimize performance without increasing cost




Frees developers from manually loading/unloading tables

Adaptive Compression: Reduces Cost and Improves Load Performance



HeatWave Lakehouse Extends Support to Semi-Structured Data

- JSON data in CSV, Parquet, and Avro file formats can be processed by HeatWave
- Support extended to newline-delimited JSON files
- Ease of parsing and streaming has made it the most popular JSON format
- NDJSON data ingestion and processing scales similarly to structured file formats

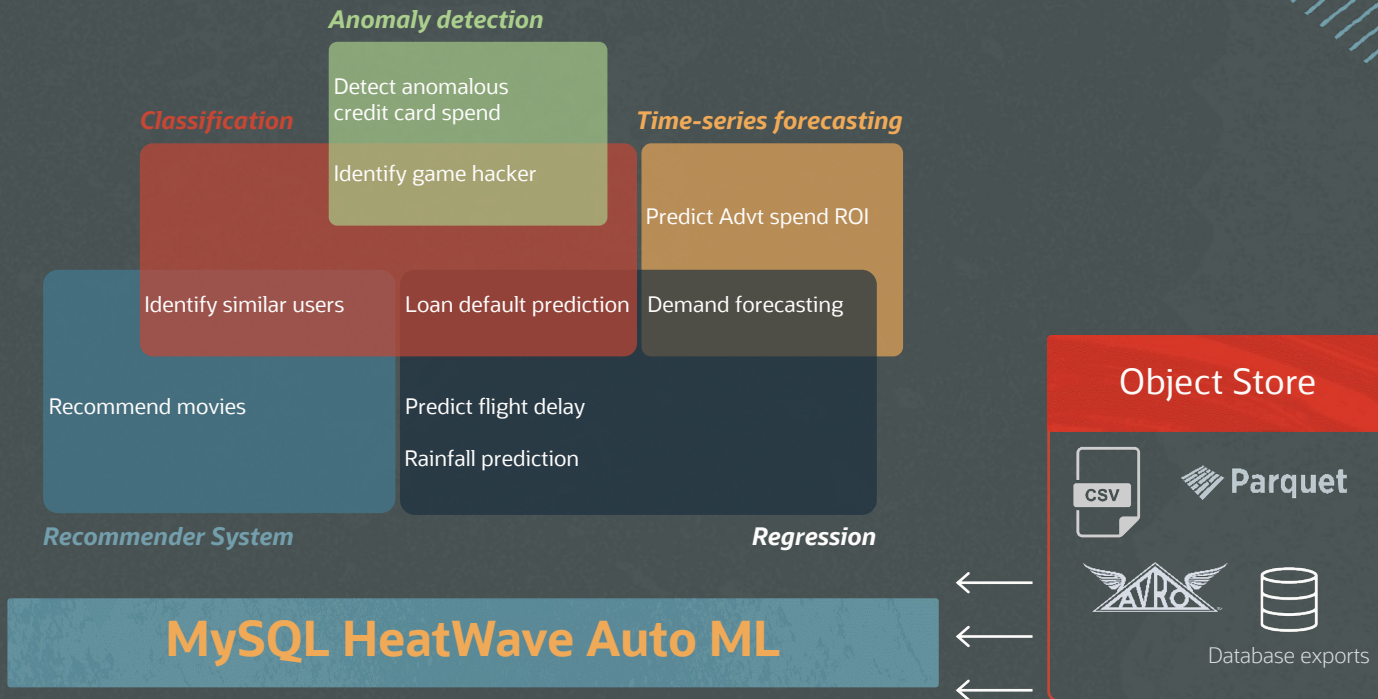


Example NDJSON file

```
...  
{ "name": "Jane", "academics": { "undergraduate": "MIT", "graduate": "UT Austin" }, "age": 24 }  
{ "name": "Jill", "academics": { "undergraduate": "Madison", "graduate": "Stanford" }, "age": 27 }  
...
```

In-DB Machine Learning

- Automated training, prediction, and explanation
- No need to ETL data to a separate ML solution
- Easy for data & business analyst
- Support for What-if scenarios
- All processing inside DB
- No additional cost
- 25x faster than Redshift



Machine Learning Use Cases

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

Rainwater level

Air pollution

Manufacturing

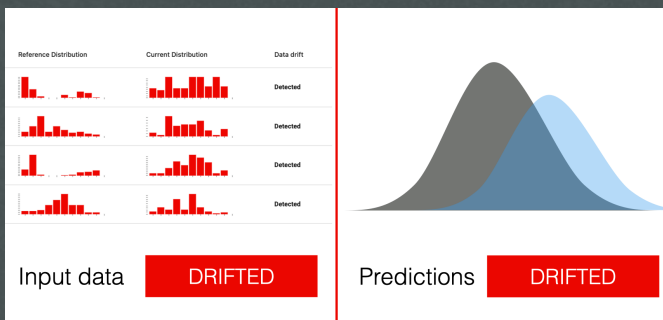
Reduce warranty claims

Defective part identification

Detect anomalies in supplies



Data Drift Monitoring



- Train a data drift detector based on Autoencoder (AE)
- Use the detector to monitor data drift in production
- The detector:
 - computes a *reconstruction error* for new incoming samples
 - updates the *cumulative drift metric*
- If the metric exceeds a threshold, automatically recommends that the model should be retrained

```
MySQL> CALL sys.ML_TRAIN('mlcorpus_v5.diamonds_train', 'price', JSON_OBJECT('task','regression'), @model);

MySQL> CALL sys.ML_MODEL_LOAD(@model, NULL);

MySQL> CALL sys.ML_PREDICT_TABLE('mlcorpus_v5.diamonds_test', @model, 'mlcorpus_v5.diamonds_predictions_experiment_results',
JSON_OBJECT('additional_details', TRUE));


MySQL> SELECT ml_results FROM diamonds_predictions_experiment_results WHERE json_extract(ml_results, '$.drift.metric') > 0.5 LIMIT 10;
+-----+
| ml_results                                                                                                                                            |
+-----+
| {"predictions": {"price": 4769.22265625}, "drift": {"metric": 0.69, "attribution_percent": {"cut": 100.0, "carat": 0.0, "clarity": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 2610.075439453125}, "drift": {"metric": 0.57, "attribution_percent": {"color": 91.25, "cut": 8.75, "carat": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 2725.368896484375}, "drift": {"metric": 0.54, "attribution_percent": {"cut": 100.0, "carat": 0.0, "clarity": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 7102.55224609375}, "drift": {"metric": 2.49, "attribution_percent": {"z": 64.53, "y": 16.86, "x": 11.58}}} |
| {"predictions": {"price": 3622.7236328125}, "drift": {"metric": 0.55, "attribution_percent": {"color": 81.2, "cut": 18.8, "carat": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 3879.93701171875}, "drift": {"metric": 2.24, "attribution_percent": {"z": 70.23, "y": 15.57, "x": 9.89}}} |
| {"predictions": {"price": 566.2338256835938}, "drift": {"metric": 0.67, "attribution_percent": {"color": 96.65, "cut": 3.35, "carat": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 2495.825439453125}, "drift": {"metric": 0.64, "attribution_percent": {"cut": 100.0, "carat": 0.0, "clarity": |
| 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 421.9180603027344}, "drift": {"metric": 0.58, "attribution_percent": {"color": 100.0, "carat": 0.0, |
| "clarity": 0.0}}}                                                                                                                                            |
| {"predictions": {"price": 325.4655456542969}, "drift": {"metric": 0.53, "attribution_percent": {"color": 100.0, "carat": 0.0, |
| "clarity": 0.0}}}                                                                                                                                            |
+-----+
+-----+
```



Anomaly Detection for MySQL Logs

NOW TRAINED FOR MYSQL LOGS

- Heatwave processes and generalizes incoming machine logs, then builds a tailored anomaly detection model
- This model helps in identifying anomalies in logs, enabling effective preventative maintenance and root cause analysis



- [2024-03-05 13:28:59 - 2024-03-05 13:29:20] Group replication-related failure
[GCS] Failure reading from fd=<:NUM:> <:*:> from <:IP:>:<:NUM:>
- [2024-03-05 13:29:25 - 2024-03-05 13:29:27] 'This server is not able to reach a majority of members in the group. This server will now block all updates. The server will remain blocked until contact with the majority is restored. It is possible to use group replication force members to force a new group membership.'
- [2024-03-05 13:40:59 - 2024-03-05 13:41:21] Potential connection leak in group replication
[GCS] Old incarnation found while trying to add node
- [2024-03-05 17:26:31] Database was not shutdown normally! Starting crash recovery. Starting to parse redo log at lsn = <:NUM:>



Server Improvements



Announcing: MySQL 8.4 LTS

MySQL Long-Term Support (LTS) Releases

- Stable: bugfix & security patches only
- Backwards compatibility
- Every 2 years
- Support lifecycle: 5y premier + 3y extended

MySQL Innovation Releases

- Leading-edge innovations
- Easy migration between LTS & Innovation
- Every quarter
- Support lifecycle: short term

MySQL 8.4 LTS

- Released: April 2024
- Supported until: April 2032 (8 years from GA)



JavaScript Stored Programs

```
CREATE FUNCTION construct_url (path VARCHAR(50),
search VARCHAR(20)) RETURNS VARCHAR(100)
LANGUAGE JAVASCRIPT AS $$
  let url = `${path}${search} &&
    !search.startsWith('?') ? '?' : ''}${search} ?? '`;
  return encodeURIComponent(url);
  $$
```

```
SELECT construct_url('/page', 'query=шел лы');
/page?query=%D1%88%D0%B5%D0%BB%D0%BB%D1%8B
```

```
CREATE PROCEDURE update_item_urls(OUT url_count INT)
LANGUAGE JAVASCRIPT AS $$
  let result = mysql.getSession().runSql(
    `UPDATE my_table
     SET url = construct_url(path, CONCAT('item=',product))
     WHERE product IS NOT NULL`
  );
  url_count = result.getAffectedItemsCount();
  $$
```

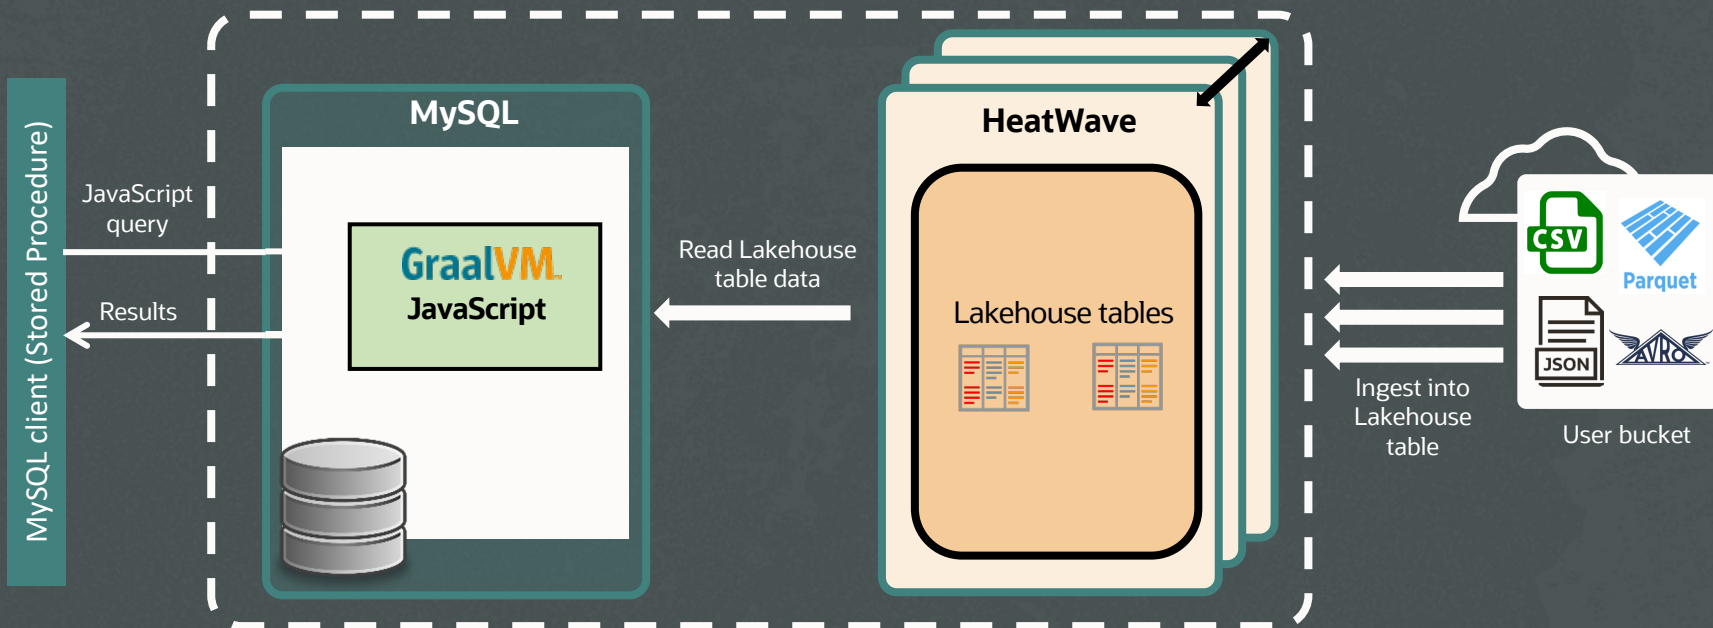
- Seamless MySQL ↔ JavaScript type conversion for input / output arguments
- First-class objects
- Can be used anywhere a SQL stored function can be used – e.g. SELECT, WHERE, ORDER BY
- Support for DML, DDL, Views
- Existing XDevAPI used to execute SQL inside JavaScript

GA for MySQL HeatWave

Currently in Beta for on-premises MySQL EE




Run JavaScript on files in Lakehouse



Performance Improvements: Thread Pool

HIGHLY SCALABLE THREAD-HANDLING MODEL

Sysbench OLTP Read/Write, "pareto" access pattern, 100GB data size



9x Better Scalability: Sysbench OLTP Read/Write



Accelerated Development



MySQL for Developers License

FREE DOWNLOAD OF MYSQL ENTERPRISE EDITION FROM OTN

Full access to MySQL Enterprise Edition

- Enterprise Server
- Backup
- Router
- Shell
- Connectors
- JavaScript

Learn, Develop, Prototype

 MySQL Enterprise Masking De-identify, Anonymize Sensitive Data	 MySQL Enterprise TDE AES 256 encryption, Key Management	 MySQL Enterprise Authentication External Authentication Modules	 MySQL Enterprise Encryption Public/Private Key Cryptography, Asymmetric Encryption
 MySQL Enterprise Firewall Block SQL Injection Attacks, Intrusion Detection	 MySQL Enterprise Audit User Activity Auditing, Regulatory Compliance	 MySQL Enterprise Thread Pool Performance & Scalability for enterprise workloads	 MySQL Enterprise Backup Secure Backups, AES 256 encryption

Download Now

<https://www.oracle.com/mysql/technologies/mysql-enterprise-edition-downloads.html>



MySQL Shell for VS Code

The screenshot shows the MySQL Shell for VS Code interface. The left sidebar displays a database tree with the following structure:

- Localhost 8.0
 - information_schema
 - performance_schema
 - sakila
 - actor
 - address
 - category
 - city
 - country
 - customer
 - film
 - film_actor
 - Columns
 - actor_id
 - film_id
 - last_update
 - Indexes
 - Foreign Keys
 - Triggers
 - film_category
 - film_text
 - Inventory
 - language
 - payment
 - rental
 - staff

The main area displays the following table:

actor_id	first_name	last_name	last_update
1	PENELOPE	GUINNESS	2021-09-28 22:18:53
2	NICK	WAHLBERG	2006-02-15 04:34:33
3	ED	CHASE	2006-02-15 04:34:33
4	JENNIFER	DAVIS	2006-02-15 04:34:33
5	JOHNNY	LOLLOBRIGIDA	2006-02-15 04:34:33
6	BETTE	NICHOLSON	2006-02-15 04:34:33
7	GRACE	NOSTEL	2006-02-15 04:34:33
8	MATTHEW	JOHNSON	2006-02-15 04:34:33
9	JOE	SWANK	2006-02-15 04:34:33
10	CHRISTIAN	GABLE	2006-02-15 04:34:33

The main area also displays a pie chart showing the distribution of languages. The chart is divided into segments for German, English, French, Italian, Mandarin, and Japanese.

Database Notebook Interface

- Write, Execute, Edit

MySQL Shell GUI Console

- Full Power of MySQL Shell

Full MySQL HeatWave Integration

- Manage MySQL HeatWave Instances



MySQL Operator for Kubernetes



Automated deployment & management

- Server
- Router
- HA/DR

Self-healing

Backup & Restore

Scaleup/Scaledown

Rolling upgrades

Configuration Management

Database Cloning

Private container registries

CNCF cert-manager support

MySQL Enterprise Edition



REST Service

FAST, SECURE HTTPS ACCESS FOR YOUR MYSQL DATA

RESTful Web Services

- Auto REST for tables, views and procedures
- {JSON} responses
- Paged results
- Developer support (GUI, CLI, API)

MySQL Shell for VS Code

- GUI frontend for MRS management
- RESTful Web Service creation
- Interactive documentation
- CLI & scripting support

Built in User Management

- Support for popular OAuth2 services
- Use Role, Group & Hierarchy Management
- User management GUI
- CLI & scripting support



Observability & Management



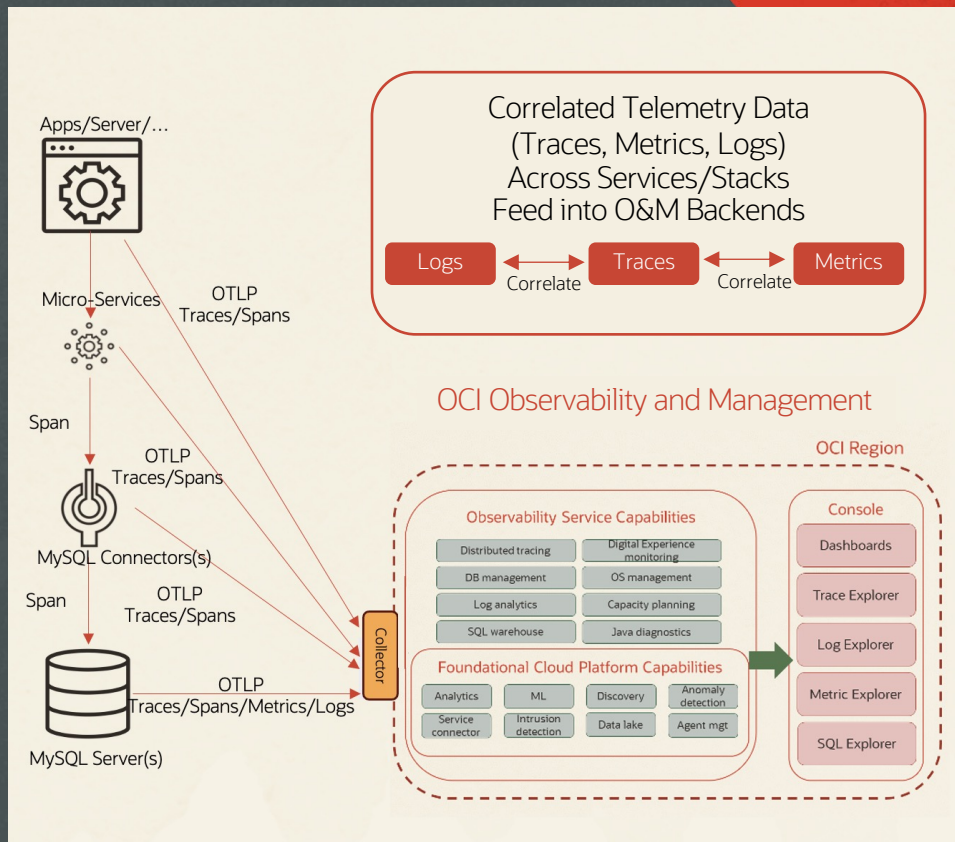
OpenTelemetry Support

OpenTelemetry (Otel)

- Open standard for telemetry data
 - Cloud Native Compute Foundation (CNCF) project
 - OCI is CNCF Platinum member
- Provides technology to collect and export telemetry
 - APIs, libraries, agents, and instrumentation

MySQL integration

- Includes Otel libraries to emit Traces, Metrics, Logs
- Traces include Spans (unit of work context)
- Metrics – choose from 400+ metrics (meters and gauges)
- Logs - in progress for 9.0



Cloud Service for Managing MySQL

MONITORING & DIAGNOSTICS*

Unified fleet monitoring and management for on-prem and cloud databases

- Monitor heterogeneous fleet (MySQL, Oracle)
- Detailed availability, performance, and configuration metrics for DevOps

Easy load and performance analysis

- Quickly identify expensive queries
- Visualize query activity for fast troubleshooting

* In GA for MySQL HW; ETA for on-premises: July 2024

PREDICTIVE INSIGHTS**

Smart capacity planning

- Forecast demand for changing workloads
- Machine learning seasonality models
- Automatic prediction of near-term issues

SQL insights

- ML driven performance diagnostic insights based on curated SQL data
- Insights provided at SQL, database, and fleet levels
- Interactive dashboards to visualize, verify, and investigate problems

** Expected availability: July 2024

Enterprise Manager for MySQL

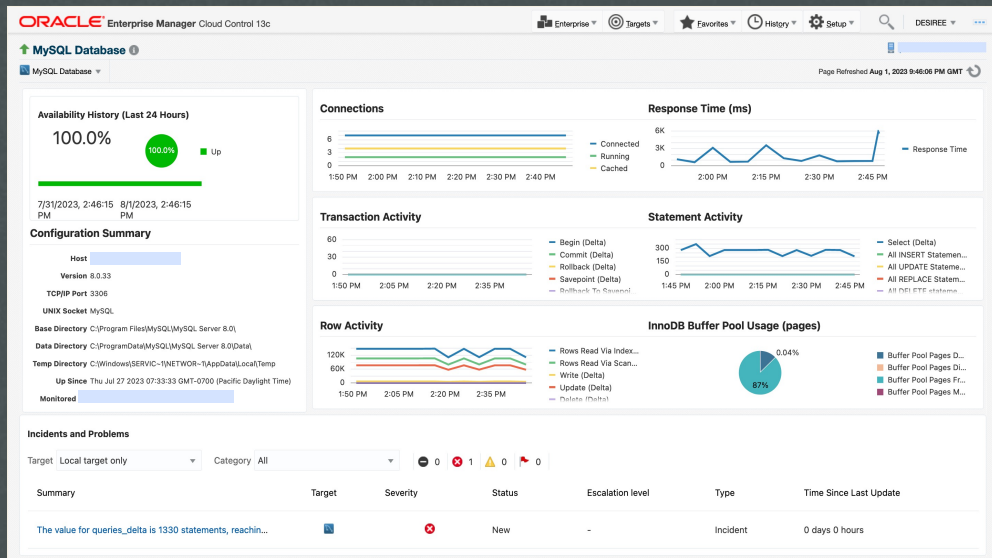
MONITORING, DIAGNOSTICS AND COMPLIANCE

Monitoring

- Status and performance monitoring
- Query Analyzer to identify slow queries
- Corrective Actions to automate alert resolution
- Metric Extensions for custom metrics
- Incident management with event compression
- Notifications: Email, SNMP traps, Webhooks, Slack
- Integrations: ServiceNow, MS SCOM, PagerDuty, Jira

Compliance

- Compliance Standards to evaluate against MySQL configuration and security best practices
- Compliance Dashboard to monitor compliance scores and violations across MySQL fleet



Innovative Organizations Run Their Business on MySQL

Social

facebook



LinkedIn



Pinterest

E-Commerce

Booking.com

NETFLIX

UBER



淘宝网
Taobao.com

阿里巴巴
Alibaba.com

SaaS



GitHub

HubSpot

zendesk



New Relic

Finance



J.P.Morgan

citi



VISA



Manufacturing

TESLA



CAT





Thank you

