

---

# MySQL HeatWave Release Notes

## Abstract

This document contains release notes for the changes in each release of MySQL HeatWave.

For additional MySQL HeatWave documentation, see [MySQL HeatWave User Guide](#).

Updates to these notes occur as new product features are added, so that everybody can follow the development process. If a recent version is listed here that you cannot find on the download page (<https://dev.mysql.com/downloads/>), the version has not yet been released.

The documentation included in source and binary distributions may not be fully up to date with respect to release note entries because integration of the documentation occurs at release build time. For the most up-to-date release notes, please refer to the online documentation instead.

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit the [MySQL Forums](#), where you can discuss your issues with other MySQL users.

Document generated on: 2025-08-04 (revision: 30411)

## Table of Contents

Preface and Legal Notices .....	2
Changes in MySQL HeatWave .....	4
Changes in MySQL HeatWave 9.4.0 (2025-07-22, General Availability) .....	4
Changes in MySQL HeatWave 9.3.2 (2025-06-24, General Availability) .....	6
Changes in MySQL HeatWave 9.3.1 (2025-05-27, General Availability) .....	7
Changes in MySQL HeatWave 9.3.0 (2025-04-15, General Availability) .....	8
Changes in MySQL HeatWave 9.2.2 (2025-03-18, General Availability) .....	9
Changes in MySQL HeatWave 9.2.1 (2025-02-20, General Availability) .....	10
Changes in MySQL HeatWave 9.2.0 (2025-01-24, General Availability) .....	11
Changes in MySQL HeatWave 9.1.2 (2024-12-17, General Availability) .....	12
Changes in MySQL HeatWave 9.1.1 (2024-11-19, General Availability) .....	13
Changes in MySQL HeatWave 9.1.0 (2024-10-15, General Availability) .....	14
Changes in MySQL HeatWave 9.0.1-u1 (2024-09-02, General Availability) .....	15
Changes in MySQL HeatWave 9.0.1 (2024-07-23, General Availability) .....	17
Changes in MySQL HeatWave 9.0.0 (2024-07-02, General Availability) .....	17
Changes in MySQL HeatWave 8.4.0 (2024-04-30, General Availability) .....	20
Changes in MySQL HeatWave 8.3.0-u2 (2024-02-26, General Availability) .....	20
Changes in MySQL HeatWave 8.3.0 (2024-01-17, General Availability) .....	21
Changes in MySQL HeatWave 8.2.0-u2 (2023-12-19, General Availability) .....	21
Changes in MySQL HeatWave 8.2.0-u1 (2023-12-05, General Availability) .....	22
Changes in MySQL HeatWave 8.2.0 (2023-10-25, General Availability) .....	23
Changes in MySQL HeatWave 8.1.0-u4 (2023-09-26, General Availability) .....	24
Changes in MySQL HeatWave 8.1.0-u3 (2023-09-12, General Availability) .....	24
Changes in MySQL HeatWave 8.1.0-u2 (2023-08-08, General Availability) .....	25
Changes in MySQL HeatWave 8.1.0-u1 (2023-07-25, General Availability) .....	25
Changes in MySQL HeatWave 8.1.0 (2023-07-18, General Availability) .....	26
Changes in MySQL HeatWave 8.0.33-u3 (2023-06-09, General Availability) .....	26
Changes in MySQL HeatWave 8.0.33 (2023-04-27, General Availability) .....	27

Changes in MySQL HeatWave 8.0.32-u2 (2023-02-21, General Availability) .....	27
Changes in MySQL HeatWave 8.0.32-u1 (2023-02-21, General Availability) .....	27
Changes in MySQL HeatWave 8.0.32 (2023-01-17, General Availability) .....	28
Changes in MySQL HeatWave 8.0.31 (2022-10-11, General Availability) .....	29
Changes in MySQL HeatWave 8.0.30-u1 (2022-09-06, General Availability) .....	30
Changes in MySQL HeatWave 8.0.30 (2022-07-26, General Availability) .....	31
Changes in MySQL HeatWave 8.0.28-u3 (2022-04-19, General Availability) .....	33
Changes in MySQL HeatWave 8.0.28-u2 (2022-03-29, General Availability) .....	33
Changes in MySQL HeatWave 8.0.28-u1 (2022-02-15, General Availability) .....	34
Changes in MySQL HeatWave 8.0.27-u3 (2021-12-15, General Availability) .....	34
Changes in MySQL HeatWave 8.0.27-u2 (2021-12-07, General Availability) .....	34
Changes in MySQL HeatWave 8.0.26-u2 (2021-09-21, General Availability) .....	35
Changes in MySQL HeatWave 8.0.26-u1 (2021-08-10, General Availability) .....	35
Changes in MySQL HeatWave 8.0.26 (2021-07-23, General Availability) .....	36
Changes in MySQL HeatWave 8.0.25 (2021-05-11, General Availability) .....	37
Changes in MySQL HeatWave 8.0.24 (2021-04-20, General Availability) .....	38
Changes in MySQL HeatWave 8.0.23-u2 (2021-03-15, General Availability) .....	39
Changes in MySQL HeatWave 8.0.23-u1 (2021-02-09, General Availability) .....	39
Index .....	40

## Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL HeatWave.

### Legal Notices

Copyright © 1997, 2025, Oracle and/or its affiliates.

#### License Restrictions

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

#### Warranty Disclaimer

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

#### Restricted Rights Notice

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including

any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

### **Hazardous Applications Notice**

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

### **Trademark Notice**

Oracle, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

### **Third-Party Content, Products, and Services Disclaimer**

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

### **Use of This Documentation**

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

## **Documentation Accessibility**

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

## Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

## Changes in MySQL HeatWave

### Changes in MySQL HeatWave 9.4.0 (2025-07-22, General Availability)



#### Note

MySQL HeatWave User Guide has been restructured and improved to provide better visibility of high-level tasks.



#### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports the recording of hyperparameters of trained models, enhancing the transparency and reproducibility of machine learning processes. This update allows for more detailed insights into model training, enabling better optimization and management of machine learning workflows. With this feature, you can access and utilize hyperparameter information to refine your models, leading to improved performance and accuracy.

For more information, see [Model Metadata](#). (WL #16946)

- MySQL HeatWave now supports a new feature called *NL2ML*, which provides an intuitive natural-language interface to machine learning (ML). With *NL2ML*, you can ask questions about MySQL HeatWave AutoML and receive step- by-step guidelines on how to use ML for specific business problems. This innovative feature allows you to generate AutoML queries that can be copied and executed, making it easier to get started with ML. Additionally, *NL2ML* offers a chat-bot interface for follow-up questions and leverages your schemas, tables, and columns to create personalized AutoML queries.

For more information, see [Learn About MySQL HeatWave AutoML Using OCI Generative AI](#). (WL #16785)

- MySQL HeatWave AutoML now supports enhanced memory management capabilities, enabling more efficient resource utilization. MySQL HeatWave can now check memory usage and release resources based on severity levels, ensuring optimal performance and reliability. This enables better management of memory between MySQL HeatWave GenAI and MySQL HeatWave AutoML. (WL #16759)
- The `rpdd_nodes` performance schema table includes new columns that enable you to monitor memory usage by MySQL HeatWave AutoML on MySQL HeatWave nodes.

For more information, see [The rpd\\_nodes Table](#). (WL #16761)

## MySQL HeatWave GenAI

- MySQL HeatWave GenAI now supports using [Dedicated Clusters from OCI Generative AI Service](#), providing you with more flexibility and capabilities to leverage the power of generative AI. With this addition, MySQL HeatWave GenAI continues to expand its feature set, offering a more comprehensive and robust solution for you to explore and innovate.



### Important

Support for `mistral-7b-instruct-v1` and `llama3-8b-instruct-v1` in MySQL HeatWave LLMs has now been deprecated.

MySQL 9.4.1 onwards, if you use any of the deprecated in-MySQL HeatWave LLMs: `llama2-7b-v1`, `mistral-7b-instruct-v1`, or `llama3-8b-instruct-v1`, then MySQL HeatWave GenAI will automatically replace it with `llama3.2-3b-instruct-v1`.

For more information, see [Supported LLMs, Embedding Models, and Languages](#). (WL #16836)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports creation and modification of Lakehouse tables using SQL syntax for setting table options, in addition to the existing way of using JSON strings for setting the `ENGINE_ATTRIBUTE`. This provides another way for you to define and load data from Object Storage for Lakehouse tables.

For more information, see [Lakehouse External Table SQL Syntax](#). (WL #16837)

- The `CREATE EXTERNAL TABLE` syntax in MySQL HeatWave Lakehouse has been enhanced to provide a more streamlined way of creating tables. The primary and secondary engines for MySQL HeatWave Lakehouse tables are now determined by session variables, eliminating the need to specify engines explicitly.

For more information, see [Lakehouse External Table Syntax](#). (WL #16831)

## MySQL HeatWave

- MySQL HeatWave now supports the use of the `URI` keyword in addition to the `URL` keyword for loading data with Bulk Load and exporting query results.

For more information, see [Bulk Ingest Data](#) and [Export Query Results to Object Storage](#). (WL #16988)

- Bulk Load now supports loading tables without primary keys. (WL #16805)
- The syntax for creating MySQL HeatWave temporary tables now uses the `ENGINE=RAPID` option instead of `SECONDARY_ENGINE=RAPID` making `RAPID` the primary engine for these tables. This enhancement also lets you specify the primary engine for MySQL HeatWave temporary tables using the `default_tmp_storage_engine` variable.

For more information, see [Create MySQL HeatWave Temporary Tables](#). (WL #16898)

- MySQL HeatWave now supports enhanced offload debuggability by introducing significant enhancements, including unified error handling and improved integration with MySQL Shell. This enhancement gives you access to more informative error messages and workarounds, making it easier to diagnose and resolve offload errors. (WL #16409)

## Changes in MySQL HeatWave 9.3.2 (2025-06-24, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports drift detection for anomaly detection, enabling you to identify changes in data distribution. With this update, you can review the model metadata and enable the `additional_details` option with the `ML_PREDICT_TABLE` and `ML_PREDICT_ROW` routines to get valuable insights into data drift, enhancing the overall anomaly detection experience.

For more information, see [Analyze Data Drift](#). (WL #16368)

### MySQL HeatWave GenAI

- Content generation using MySQL HeatWave GenAI now supports speculative decoding, which enables faster response token generation and speeds up text generation if the target Large Language Model (LLM) supports speculative decoding.

For more information, see [ML\\_GENERATE](#), [ML\\_GENERATE\\_TABLE](#), and [@chat\\_options Parameters](#). (WL #16809)

- Document ingestion using Auto Parallel Load has been enhanced to support customised segmentation of text during vector store creation. This enhancement lets you have greater control over how text is segmented, enabling more precise and tailored text analysis capabilities.

For more information, see [Ingest Files Using Auto Parallel Load](#). (WL #16683)

- MySQL HeatWave GenAI now supports automated discovery and listing of all available LLMs and embedding models in MySQL HeatWave. This enhancement lets you stay up-to-date on model availability, and enables you to make informed usage decisions especially in the case of OCI Generative AI Service models as the list of models available in MySQL HeatWave might change before a new version of MySQL is available.

For more information, see [Supported LLMs, Embedding Models, and Languages](#). (WL #16664)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports the direct loading of compressed files from Object Storage into tables, enhancing data ingestion capabilities and simplifying data pipelines. This update lets you load `gzip`, `zip`, and `bzip2` compressed files that contain CSV and JSON data into MySQL HeatWave. This enhancement streamlines the process of working with compressed data, making it more efficient and cost effective for you to manage and analyze your data within MySQL HeatWave Lakehouse.

For more information, see [Lakehouse External Table Syntax](#). (WL #16650)

- MySQL HeatWave Lakehouse now supports loading vector columns from CSV and Parquet files, enabling you to bring your own custom vector embeddings into MySQL HeatWave and leverage MySQL

HeatWave GenAI features. With this enhancement, you can optimize and manage semantic search capabilities for your datasets and integrate with the large ecosystem of GenAI capabilities.

For more information, see [Supported File Formats and Data Types](#). (WL #16563)

## Changes in MySQL HeatWave 9.3.1 (2025-05-27, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

### MySQL HeatWave GenAI

- MySQL HeatWave GenAI is now available on all MySQL HeatWave Cluster shapes, including the [HeatWave.Free](#) shape.

Also, the following in-database LLMs are now available in MySQL HeatWave GenAI: [llama3.2-3b-instruct-v1](#), [llama3.2-1b-instruct-v1](#), [llama3.1-8b-instruct-v1](#), and [mistral-7b-instruct-v3](#). With the introduction of new LLMs, MySQL HeatWave GenAI now uses [llama3.2-3b-instruct-v1](#) as the default LLM for content generation instead of [mistral-7b-instruct-v1](#).

For more information, see [Supported LLMs, Embedding Models, and Languages](#). (Bug #37551135, WL #16509)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports seamless loading of highly-encoded Parquet files, including those with sparse data, highly repetitive values, and consecutive values. (WL #16812)
- MySQL HeatWave Lakehouse now supports unified loading of files using different URI formats. With this enhancement, you can load files from Object Storage using any of the three URI formats: OCIFS URI, PAR URI, and native URI. These URIs support patterns, prefixes, or names, which makes it easier to work with large datasets. Additionally, error and warning messages have been improved to include the relevant URI information for better troubleshooting.

For more information, see [Access Object Storage with Uniform Resource Identifiers \(URI\)](#). (WL #16822)

- MySQL HeatWave now supports exporting query results to Object Storage as Newline Delimited-JSON (ND-JSON) files. Using ND-JSON files lets you export query results as line-separated JSON logs, which enables easy and efficient integrations of your data with other web applications. This enhancement also makes it easier to load the exported results into a table with a single column of [JSON](#) data type, or use the [HEATWAVE\\_LOAD](#) routine to load the exported results into MySQL HeatWave Cluster for further processing.

For more information, see [Export Query Results to ND-JSON Files](#). (WL #16778)

- MySQL HeatWave now supports [MEDIUMTEXT](#), [LONGTEXT](#), and [JSON](#) data types with a column width of up to 4MB for Lakehouse tables. With this enhancement, you can leverage the expanded capabilities of MySQL HeatWave Lakehouse to load and manage larger datasets, boosting overall performance and productivity.

For more information, see [Lakehouse Limitations for all File Formats](#). (WL #16472)

## MySQL HeatWave

- MySQL HeatWave now supports automated partition management, extending the Auto Parallel Load feature of automatically loading and unloading only the necessary partitions in partitioned tables based on usage. This enhancement accelerates workloads on partitioned tables while offering memory savings, allowing for more efficient data management and improved performance.

For more information, see [Automatic Loading and Unloading of DB System Tables and Partitions](#). (WL #15994)

- MySQL HeatWave now supports creating asynchronous tasks for long-running queries and commands, allowing them to run in the background as separate tasks. This enhancement enables you to run complex queries without blocking other operations, thus enhancing the overall performance and usability of MySQL HeatWave.

For more information, see [Run Tasks Asynchronously](#). (WL #16539)

## Changes in MySQL HeatWave 9.3.0 (2025-04-15, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.



### Important

As of MySQL 9.3.0, to help you generate better quality embeddings, MySQL HeatWave GenAI uses `multilingual-e5-small` as the default embedding model for encoding documents in all supported languages including English. This means that `minilm` is no longer used as the default embedding model for encoding English documents.

This default value change impacts the following processes:

- New vector store tables are created using `multilingual-e5-small` by default.
- Retrieval augmented generation (RAG) searches tables created using the default embedding model, `multilingual-e5-small`, unless you explicitly specify the embedding model to use.

- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

## MySQL HeatWave GenAI

- MySQL HeatWave GenAI now supports partial failures with better error reporting in batch queries that are run using the `ML_GENERATE_TABLE`, `ML_RAG_TABLE`, and `ML_EMBED_TABLE` routines. This enhancement allows queries on certain rows to fail in case of errors without discarding successful work on other rows. (WL #16749)

## MySQL HeatWave Lakehouse

- MySQL HeatWave now supports automatic recovery of loaded Lakehouse tables from the MySQL HeatWave Storage Layer after a planned or unplanned restart of the primary DB system. When a DB system restarts, data for Lakehouse tables is loaded from the [MySQL HeatWave Storage Layer in Object Storage](#). After a successful load, the data is current as of the last refresh of the table. If this operation fails, data for these Lakehouse tables is then loaded from the Object Storage bucket.

This feature extends the existing functionality of automatic recovery from Storage Layer for Standalone DB systems to High Availability DB systems, thus reducing the recovery time, increasing the availability, and improving the overall system reliability.

For more information, see [About MySQL HeatWave](#). (WL #16758)

## MySQL HeatWave

- MySQL HeatWave now supports the creation of temporary tables, which are stored in a hybrid columnar format within the MySQL HeatWave Cluster. If you use temporary tables to store intermediate results during data transformation, aggregation or consolidation for data analysis or reporting purposes, you can now accelerate the processes with MySQL HeatWave.

For more information, see [Create MySQL HeatWave Temporary Tables](#). (WL #16541)

- Bulk Load, which is used for ingesting data into the DB System, now supports tables that include the [VECTOR](#) data type. This enables you to easily and quickly import existing or pre-generated embeddings to MySQL HeatWave for your MySQL HeatWave GenAI workloads.

For more information, see [Bulk Ingest Data Type Support](#) (WL #16510)

- MySQL HeatWave now supports the use of the [ANALYZE TABLE](#) statement to update statistics for tables loaded into the MySQL HeatWave Cluster. This enhancement improves query performance by refreshing critical statistics used during query processing. In mixed or Lakehouse workloads, where changing data characteristics can degrade performance, this feature helps maintain optimal query execution by keeping statistics up-to-date.

For more information, see [Analyze Tables](#). (WL #16641)

- MySQL HeatWave now supports acceleration of queries with general quantified comparison predicates.

For more information, see [Optimizer Notes](#) and [Optimizing ANY and ALL Subqueries](#). (WL #13052)

## Changes in MySQL HeatWave 9.2.2 (2025-03-18, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.



### Important

MySQL 9.3.0 onwards, to help you generate better quality embeddings, MySQL HeatWave GenAI will use [multilingual-e5-small](#) as the default embedding model for encoding documents in all supported languages including English. This means that [minilm](#) will not be used as the default embedding model for encoding English documents.

This default value change will impact the following processes:

- New vector store tables will be created using `multilingual-e5-small` by default.
- Retrieval augmented generation (RAG) will search tables created using the default embedding model, `multilingual-e5-small`, unless you explicitly specify the embedding model to use.

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

## MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports log anomaly detection, a new task that enables you to identify unusual patterns in log data. This feature is designed to help you detect and investigate potential issues more efficiently.

For more information, see [Anomaly Detection for Logs](#). (WL #16579)

## MySQL HeatWave Lakehouse

- A new table state, `RECOVERYFAILED_RPDGSTABSTATE`, has been introduced for Lakehouse tables. During a MySQL HeatWave Cluster restart, if a Lakehouse table cannot be recovered, its state is set to `RECOVERYFAILED_RPDGSTABSTATE` to indicate that it could not be recovered. The cluster remains usable even if recovery of Lakehouse tables fail. Tables in the `RECOVERYFAILED_RPDGSTABSTATE` state cannot be used for queries or change propagation. You must reload tables in this state by running `SECONDARY_UNLOAD` followed by `SECONDARY_LOAD`, which lets you identify errors and take corrective action.

For more information, see [The rpd\\_tables Table](#). (WL #16626)

- Lakehouse now supports detailed warnings summary for load and inference operations. The summary is included in the warning messages, providing an accurate list of warning codes and their respective counts. This enhancement helps you better understand and manage warnings encountered during Lakehouse operations.

For more information, see [MySQL HeatWave Lakehouse Error Messages](#). (WL #16637)

## MySQL HeatWave

- MySQL HeatWave now supports columns wider than 64 KB, with a new maximum limit of 4 MB.

For more information, see [Column Limits](#). (WL #16289)

## Changes in MySQL HeatWave 9.2.1 (2025-02-20, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave](#)

## MySQL HeatWave GenAI

- Context search in MySQL HeatWave GenAI now lets you use your own embedding tables. This enhancement enables more flexible query enrichment for Large Language Models (LLMs). This update extends the search capabilities of the [ML\\_RAG](#), [ML\\_RAG\\_TABLE](#), and [HEATWAVE\\_CHAT](#) routines to include your own embedding tables.

For more information, see [Use Your Own Embeddings](#) and [Run MySQL HeatWave Chat](#). (WL #16537)

## MySQL HeatWave

- MySQL HeatWave now supports automatic recovery of loaded InnoDB tables from Storage Layer after a planned or unplanned restart of the primary DB system. This feature extends the existing functionality of automatic recovery from Storage Layer for Standalone DB systems to HA DB systems, thus reducing the recovery time, increasing the availability, and improving the overall system reliability. However, Lakehouse tables in HA DB systems still need to be recovered from Object Storage.

For more information, see [About MySQL HeatWave](#). (WL #16627)

- MySQL HeatWave now supports automated zone map construction on non-primary key columns, improving RAPID query processing performance. This feature uses a workload-driven approach to select the most beneficial zone map columns based on query history, continuously adapting to the workload.

For more information, see [About MySQL HeatWave](#). (WL #15993)

- The [POOL\\_TYPE](#) column in MySQL HeatWave Performance Schema tables no longer includes the [RAPID\\_LOAD\\_POOL\\_](#) prefix. The possible values are now [SNAPSHOT](#) and [TRANSACTIONAL](#).

For more information, see [The rpd\\_tables Table](#). (WL #16630)

- For tracking the usage of various components, MySQL HeatWave Option tracking now supports a single, fast counter, [usedCounter](#), instead of a boolean value, [used](#). The counter indicates the number of times a component is used.

For more information, see [Option Tracker](#). (WL #16721, WL #16732)

## Changes in MySQL HeatWave 9.2.0 (2025-01-24, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

## MySQL HeatWave GenAI

- The MySQL HeatWave GenAI Javascript API now supports batch processing, making it possible to perform multiple operations in a single call. This update introduces batch processing functionality to the [LLM](#) class and API convenience methods (see [Convenience Methods](#)), allowing for more efficient embedding, text-based content generation, and Retrieval Augmented Generation (RAG).

For more information, see [JavaScript GenAI API](#). (WL #16514)

## MySQL HeatWave Lakehouse

- Lakehouse Autopilot can now automatically detect the dialect for CSV files, which lets you load the CSV files without knowing the exact values of the `dialect` delimiter parameters in advance. This feature automates the detection of record delimiters, field delimiters, and composite record delimiters with field delimiters as prefixes.

For more information, see [Lakehouse External Table Syntax](#). (WL #16546)

## MySQL HeatWave

- Bulk load, which is used for ingesting data into MySQL Server, now supports the following:
  - Secondary indexes, which enable faster data loading and improved performance. This feature allows for the creation of secondary indexes during the bulk load process, reducing the need for additional `ALTER TABLE` statements.
  - Faster and more efficient bulk loading of CSV files. This enhancement allows for parallelized CSV parsing, reducing the need for redundant reads and improving overall performance.

For more information, see [Bulk Ingest Data](#). (WL #16266, WL #15986)

- MySQL HeatWave now supports JSON tables, enabling you to extract data from JSON documents and return it as relational tables.

For more information, see [JSON Functions](#). (WL #16450)

- MySQL HeatWave now supports offloading the `INSERT INTO SELECT` queries to the secondary engine, even when the table being inserted into has triggers defined on it. This enables faster query execution and improved performance for workloads that involve inserting data into tables with triggers.

For more information, see [INSERT ... SELECT Statements](#). (WL #16235)

- MySQL HeatWave now reloads data from Storage Layer after a planned or unplanned restart of Standalone DB system, thus reducing the recovery time and improving overall system reliability.

For more information, see [About MySQL HeatWave](#). (WL #14953)

## Changes in MySQL HeatWave 9.1.2 (2024-12-17, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)

## MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports more efficient memory usage, reducing the footprint of its processes during training and explanations for large datasets. (WL #16606)

## MySQL HeatWave GenAI

- MySQL HeatWave GenAI supports enhanced unstructured document ingestion with [Optical Character Recognition \(OCR\)](#) now enabled by default. (WL #16584)
- The `ML_RAG` and `ML_RAG_TABLE` routines now supports advanced retrieval options, enabling more precise and context-aware results. Three new parameters have been added to the routines:
  - `max_distance` for filtering segments based on distance from the input query.
  - `percentage_distance` for adaptive distance filtering of segments from the input query.
  - `segment_overlap` for retrieving segments adjacent to the segment nearest to the input query for providing continuous context.

These new parameters are integrated into the existing `@options` parameter, and can be enabled through a new JSON object parameter, `@retrieval_options`. (WL #16543)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports selective load, which lets you update the file paths comprising a Lakehouse table by manipulating Lakehouse metadata. The changes are incrementally applied to the table data. This feature enables you to add or remove Object Storage bucket data directories from a Lakehouse table without reloading the entire table. Changes are executed transactionally, ensuring data consistency and query isolation.

For more information, see [Add or Remove Files Using Selective Load](#). (WL #16531)

- MySQL HeatWave Lakehouse now supports enhanced error and warning messages, providing more data points for easier debugging. This update adds row numbers and character offsets to relevant warning messages in non-strict mode during secondary load, as well as updates byte offset text messages to point to column positions. JSON datatype error and warning messages have also been improved with error and warning offset information within the column.

For more information, see [MySQL HeatWave Lakehouse Error Messages](#). (WL #16530)

- MySQL HeatWave Lakehouse now supports efficient loading of Parquet files with large row groups (larger than 500MB and upto 10GB) using large shapes. This helps in overcoming previous memory budget constraints and significantly reducing load times. This enhancement also introduces row group sharing, caching, and synchronization mechanisms to minimize network requests, decoding time, and memory usage during transformation.

For more information, see [Lakehouse Limitations for the Parquet File Format](#). (WL #16512)

## Changes in MySQL HeatWave 9.1.1 (2024-11-19, General Availability)



### Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)

## MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports the Orbit forecasting model, which uses Local Global Trend (LGT). This addition expands the scope of predictive analytics and enables more advanced time-series forecasting for businesses.

For more information, see [Generate Forecasts](#). (WL #16231)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports exporting query results directly to an object store in OCI or AWS. This lets you store, transform, and persist data in CSV and Parquet formats. The new syntaxes introduced make it easier to export data, especially in command-line environments, and provide more control over the output format. This enhancement enables efficient data storage, retrieval, and transformation, improving data management and analysis workflows.

For more information, see [Export Query Results to Object Storage](#). (WL #16214)

## Changes in MySQL HeatWave 9.1.0 (2024-10-15, General Availability)

- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

## MySQL HeatWave GenAI

- MySQL HeatWave GenAI now lets you extract and encode text from images stored in unstructured documents using Optical Character Recognition (OCR). The extracted text is converted into vector embeddings and stored in a vector store the same way regular text in unstructured documents is encoded and stored in a vector store.

For more information, see [About Optical Character Recognition](#). (WL #16474)

- MySQL JavaScript Stored Programs now include a new GenAI API that you can use to call different [MySQL HeatWave GenAI routines](#) using JavaScript functions.

For more information, see [JavaScript GenAI API](#). (WL #16487)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports zone maps, which helps you improve range queries in OLAP and mixed workloads.

MySQL HeatWave uses statistics based on minimum and maximum values to create zone maps for every primary key column. It then uses the zone maps for range and point queries to only scan data chunks that are relevant for the query, and accelerates these queries by an order of magnitude. (WL #16414)

## MySQL HeatWave

- Auto Parallel Load now uses Autopilot to collect statistics about frequently used InnoDB and MySQL HeatWave tables. Auto Parallel Load then automatically loads these tables to MySQL HeatWave. If a manually loaded table might cause a memory conflict, Auto Unload unloads automatically loaded tables to free up memory. (WL #15491)

- MySQL HeatWave now supports InnoDB partitions, and `INFORMATION_SCHEMA.PARTITIONS` includes a `SECONDARY_LOAD` column that indicates whether a partition has been loaded.

When an unload operation successfully unloads the last partition, it unloads the whole table.

MySQL HeatWave now persists partition information, and during a restart, only the previously loaded partitions are loaded.

For more information, see [Load Partitions](#) (WL #15517, WL #16268)

## Changes in MySQL HeatWave 9.0.1-u1 (2024-09-02, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports semi-supervised learning for anomaly detection. This type of machine learning algorithm uses a specific set of labeled data along with unlabeled data to detect anomalies. (WL #16373)
- MySQL HeatWave AutoML now supports topic modeling. This is an unsupervised machine learning technique that is capable of scanning a set of documents, detecting word and phrase patterns within them, and automatically clustering word groups and similar expressions that best characterize the documents. (WL #16275)

### MySQL HeatWave GenAI

- MySQL HeatWave GenAI now supports multiple languages, which lets you use the MySQL HeatWave GenAI APIs in non-English languages. This includes ingestion of documents written in languages other than English, vector and similarity search, and querying these documents by giving prompts in the same language.

For more information, see [Languages](#). (WL #16475)

- MySQL HeatWave GenAI now supports large language model (LLM) inference batch processing, which lets you run the MySQL HeatWave GenAI routines on multiple queries, in parallel, across different nodes in the MySQL HeatWave Cluster. Thus, it improves the LLM inference performance and throughput while keeping the inference latency of every request unchanged.

For more information, see [ML\\_GENERATE\\_TABLE](#), [ML\\_RAG\\_TABLE](#), and [ML\\_EMBED\\_TABLE](#). (WL #16399)

- The MySQL HeatWave GenAI `ML_RAG` routine now includes new filtering options, `exclude_vector_store` and `exclude_document_name`, which let you exclude specific vector store tables or documents from context retrieval. (WL #16401)
- You can now use incremental load to refresh vector store tables. For more information, see [Load Data Incrementally into the Vector Store Table](#). (WL #15817)
- Auto Parallel Load for unstructured data provides a new `format` option setting, `auto_unstructured`, which lets you ingest multiple files with different unstructured data formats in a single load.

Additionally, for ingesting Microsoft PowerPoint and Microsoft Word files into the vector store, Auto Parallel Load introduces the following new format aliases: `pptx` and `docx`.

For more information, see [Ingest Files Using Auto Parallel Load](#). (WL #16433)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports `timestamp_format` as a `dialect` parameter to customize the format for columns of the `DATETIME` and `TIMESTAMP` data types. You can also set the formats for specific columns by using the `ENGINE_ATTRIBUTE` option. Formats set at the column level override formats set with the `dialect` parameter. For more information, see [Lakehouse External Table Syntax](#). (WL #16425)
- Some MySQL HeatWave Lakehouse error messages include the URL to the external object. The URL can contain up to 1024 characters, and exceed the error message limit of 512 characters. The truncated error message will now include a MySQL command to access the full error message.

It is now possible to filter out MySQL HeatWave Lakehouse warning messages. The console **Total Warnings** will include both the displayed and the filtered warnings count. See: [MySQL HeatWave Lakehouse Error Messages](#).

MySQL HeatWave Lakehouse now supports `max_error_count`. (WL #16141)

## MySQL HeatWave

- `GROUP_CONCAT()` function now supports the `CUBE`, `ROLLUP`, and `WITH ROLLUP` options.
  - `COUNT(DISTINCT)` function now supports the `CUBE`, `ROLLUP`, and `WITH ROLLUP` options. (WL #16365)
- MySQL HeatWave now supports the following JSON functions:
  - `JSON_INSERT()`
  - `JSON_REMOVE()`
  - `JSON_REPLACE()`
  - `JSON_SET()`
  - `JSON_ARRAY_APPEND()`
  - `JSON_ARRAY_INSERT()`
  - `JSON_MERGE_PATCH`
  - `JSON_MERGE_PRESERVE`
  - `JSON_MERGE()`
  - `JSON_SCHEMA_VALID()`
  - `JSON_SCHEMA_VALIDATION_REPORT()`

(WL #16345)

## Changes in MySQL HeatWave 9.0.1 (2024-07-23, General Availability)

This release does not include any new features.

## Changes in MySQL HeatWave 9.0.0 (2024-07-02, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Autopilot](#)
- [MySQL HeatWave GenAI](#)
- [MySQL HeatWave Lakehouse](#)
- [MySQL HeatWave](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML has increased the maximum model size from 900MB. The MySQL HeatWave cluster shape you set for a DB system now defines the total memory available to train a model and for all loaded models. For imported models, we recommend individual models have a size of 4GB or less. (WL #15931)

### MySQL HeatWave Autopilot

- MySQL HeatWave Autopilot includes improvements to cardinality estimation for queries that run in MySQL HeatWave and MySQL HeatWave Lakehouse. This improves the query plan, and provides improved performance (WL #15995)
- MySQL HeatWave Autopilot now includes Auto Indexing Advisor that offers indexing suggestions to improve the performance and latency of OLTP workloads running on the MySQL HeatWave primary engine. (WL #15663)

### MySQL HeatWave GenAI

- MySQL HeatWave GenAI introduces similarity search capabilities that support the `DISTANCE` function that measures the distance between two vector values. It can include an optional `distance_metric` that can have a value of `DOT`, `COSINE` or `EUCLIDIAN`. This is available in the primary and secondary engines. (WL #16081)
- MySQL HeatWave GenAI introduces the new native `VECTOR` data type and in-database vector store support. This provides a fast, end-to-end, fully integrated pipeline which automates the vector store creation:
  - Reading unstructured data in PDF, HTML, TXT, PPT or DOCX format from object store.
  - Parsing the text in the documents.
  - Partitioning the text into smaller paragraphs and segments.
  - Encoding the paragraphs.
  - Storing the encoded paragraphs in a standard MySQL table in MySQL HeatWave.

(WL #16106)

- MySQL HeatWave GenAI introduces a set of intuitive chat functionalities that enable enterprises to quickly build their own AI powered chatbot. It also introduces an out-of-the-box chatbot – MySQL HeatWave Chat, which integrates the chat functionalities that enable users to interact with MySQL HeatWave using natural language in addition to SQL. (WL #16248)
- MySQL HeatWave GenAI introduces in-database LLM with support for the following LLMs, see [Perform AI-Powered Search and Content Generation](#):
  - Mistral-7B-Instruct.
  - Llama-2-8B-Instruct.
  - LLMs from OCI GenAI.

(WL #15844, WL #16145)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports unstructured documents, and loads them to the GenAI vector store. Supported formats include PDF, HTML, TXT, PPT and DOCX. (WL #16100)
- MySQL HeatWave Lakehouse introduces Lakehouse Incremental Load that uses change propagation to refresh data views following changes to a loaded table. (WL #15817)

## MySQL HeatWave

- MySQL HeatWave includes improvements to query optimization that lead to better performance for low latency queries in MySQL HeatWave and MySQL HeatWave Lakehouse. (WL #16222)
- MySQL HeatWave includes improvements to change propagation performance. (WL #16148)
- MySQL HeatWave improves the debugging of offloaded queries. (WL #15988)
- MySQL HeatWave now supports the use of cursors from stored procedures. (WL #16142)
- MySQL HeatWave now has support for the [JSON](#) data type with the following functions:
  - [Cast Functions and Operators](#).
  - [COALESCE\(\)](#) and [IN\(\)](#), see: [Comparison Functions and Operators](#).
  - [Control Flow Functions and Operators](#).
  - [Mathematical Functions](#).
  - [String Functions and Operators](#).

(WL #16048)

- MySQL HeatWave now supports the following JSON functions:
  - [JSON\\_ARRAYAGG\(\)](#)
  - [JSON\\_CONTAINS\(\)](#)
  - [JSON\\_CONTAINS\\_PATH\(\)](#)
  - [JSON\\_KEYS\(\)](#)

- `JSON_OBJECTAGG()`
- `JSON_OVERLAPS()`
- `JSON_PRETTY()`
- `JSON_QUOTE()`
- `JSON_SEARCH()`
- `JSON_STORAGE_FREE()`
- `JSON_STORAGE_SIZE()`
- `JSON_TYPE()`
- `JSON_VALID()`
- `JSON_VALUE()`
- `MEMBER OF()`

(WL #16046, WL #16047)

- MySQL HeatWave now includes support for `GROUP_CONCAT()` with `DISTINCT` optimization and the `ORDER BY` clause. (WL #15906)
- MySQL HeatWave now supports dynamic query offload which uses machine learning based on query properties to determine the best engine, InnoDB or MySQL HeatWave analytic engine, to process the query. (WL #15829)
- MySQL HeatWave now supports concurrent query execution that reduces query latencies, and increases resource utilization. (WL #14826)

## Functionality Added or Changed

- MySQL InnoDB on MySQL HeatWave includes a bulk load extension to the `LOAD DATA` statement. This now uses multiple threads to load a series of CSV files to improve performance, and allow the `LOAD` operation to scale with multiple CPUs. (WL #15611)
- MySQL InnoDB on MySQL HeatWave includes a bulk load extension to the `LOAD DATA` statement. This now supports large objects with `VARCHAR` and `VARBINARY`. It also supports the following data types:
  - `BIT`.
  - `ENUM`.
  - `JSON`.
  - `SET`.
  - `TIMESTAMP`.
  - `YEAR`.
  - `TINYBLOB`, `BLOB`, `MEDIUMBLOB`, and `LONGBLOB`.
  - `TINYTEXT`, `TEXT`, `MEDIUMTEXT`, and `LONGTEXT`.

- [GEOMETRY](#), [GEOMETRYCOLLECTION](#), [POINT](#), [MULTIPOINT](#), [LINESTRING](#), [MULTILINESTRING](#), [POLYGON](#), and [MULTIPOLYGON](#),

See: [Bulk Ingest Data](#). (WL #16264, WL #16265)

## Changes in MySQL HeatWave 8.4.0 (2024-04-30, General Availability)

- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports primary key and unique key constraint validation with the [check\\_constraints](#) parameter. When you set this parameter to [true](#), a validation for primary key and unique key constraints occurs only during the initial load of a table. (WL #16094)
- MySQL HeatWave Lakehouse now supports the Newline Delimited JSON file format, see [Supported File Formats and Data Types](#). (WL #16079)
- MySQL HeatWave Lakehouse now has improved support for missing files. Auto Parallel Load now has a [validation](#) mode to validate files before loading them. (WL #16078)

### Functionality Added or Changed

- MySQL HeatWave now supports named time zones such as [MET](#) or [Europe/Amsterdam](#). (WL #16098)
- MySQL HeatWave Auto Parallel Load and Auto Unload now support a table inclusion list, as well as a table exclusion list. (WL #16070)
- MySQL InnoDB on MySQL HeatWave includes a bulk load extension to the [LOAD DATA](#) statement. This can now load MySQL Shell dump files and compressed files. There is also a progress monitor. See: [Bulk Ingest Data](#). (WL #15814, WL #15607, WL #15608, WL #15609)

## Changes in MySQL HeatWave 8.3.0-u2 (2024-02-26, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML adds prediction intervals to the forecasting task. (WL #16072)
- MySQL HeatWave AutoML adds the PCA and GLOF algorithms to the anomaly detection task. (WL #16104)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports the JSON data type. (WL #16037)
- MySQL HeatWave Lakehouse now supports Guided Load to perform pre-load validation checks and schema inference. (WL #16004)
- MySQL HeatWave Lakehouse now supports the following:

- Point-in-time-recovery.
- High Availability.
- Read Replication.
- Outbound Replication.

(WL #15926)

- MySQL HeatWave Lakehouse now supports Auto Data Compression. This dynamically determines which compression algorithm to use for each column based on its data characteristics. This provides the optimum balance between memory usage and query performance. See: [About MySQL HeatWave](#). (WL #16061)

### Functionality Added or Changed

- MySQL HeatWave primary and secondary engines now support the HyperLogLog, `HLL( )` aggregate function. This is similar to `COUNT(DISTINCT)`, but faster and with user defined precision. This is only available in MySQL HeatWave, see: [Aggregate Functions](#). (WL #15992)
- MySQL HeatWave can now reload all tables, and unload all tables. (WL #16023)
- MySQL HeatWave secondary engine now supports the `TABLESAMPLE` clause to use with `SELECT` statements. (WL #15902)

## Changes in MySQL HeatWave 8.3.0 (2024-01-17, General Availability)

- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse extends Auto Parallel Load support for external tables that have a pre-defined schema. It is now able to reconcile the inferred schema with the pre-defined table definition when there is a column type mismatch between the two. Schema inference adjusts this schema to prevent possible errors during data loading. See: [Use Lakehouse Auto Parallel Load with external\\_tables Option](#). (WL #16040)
- MySQL HeatWave Lakehouse now has improved error reporting. See: [MySQL HeatWave Lakehouse Error Messages](#). (WL #15951)

### Functionality Added or Changed

- `GREATEST( )` and `LEAST( )` comparison and string functions now support the `YEAR` data type. (WL #15910)

## Changes in MySQL HeatWave 8.2.0-u2 (2023-12-19, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)

### MySQL HeatWave AutoML

- The MySQL HeatWave AutoML recommendation task now supports content based recommendation models. See [Generate Recommendations](#). (WL #15726)

- MySQL HeatWave AutoML introduces data drift detection for classification and regression models. (WL #15866)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now supports all queries that MySQL HeatWave supports. See: [MySQL HeatWave Lakehouse Limitations](#). (WL #15730)

## Changes in MySQL HeatWave 8.2.0-u1 (2023-12-05, General Availability)

- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now uses Auto Parallel Load to load Avro files. (WL #15778)
- MySQL HeatWave Lakehouse schema inference now includes support for unicode column headers. (WL #16011)

## Functionality Added or Changed

- MySQL HeatWave secondary engine now supports the [QUALIFY](#) clause to use with [SELECT](#) statements. (WL #15864)
- Auto Data Compression dynamically determines which compression algorithm to use for each column based on its data characteristics. This provides the optimum balance between memory usage and query performance. See: [About MySQL HeatWave](#). (WL #15248)
- MySQL HeatWave now supports the following aggregate functions used as window functions:
  - [STD \( \)](#)
  - [STDDEV \( \)](#)
  - [STDDEV\\_POP \( \)](#)
  - [STDDEV\\_SAMP \( \)](#)
  - [VAR\\_POP \( \)](#)
  - [VAR\\_SAMP \( \)](#)
  - [VARIANCE \( \)](#)(WL #15367)
- MySQL HeatWave secondary engine now supports the HyperLogLog, [HLL \( \)](#) aggregate function. This is similar to [COUNT \(DISTINCT\)](#), but faster and with user defined precision. This is only available in MySQL HeatWave, see: [Aggregate Functions](#). (WL #15914)
- The [GROUP BY](#) clause permits a [CUBE](#) modifier. This is only available in MySQL HeatWave, see: [GROUP BY Modifiers](#). (WL #15843)
- MySQL HeatWave now supports the following JSON functions:
  - [->](#)

- `->>`
- `JSON_ARRAY( )`
- `JSON_DEPTH( )`
- `JSON_EXTRACT( )`
- `JSON_LENGTH( )`
- `JSON_OBJECT( )`
- `JSON_UNQUOTE( )`

See: [JSON Functions](#). (WL #15355)

## Changes in MySQL HeatWave 8.2.0 (2023-10-25, General Availability)

### Functionality Added or Changed

- MySQL HeatWave Guided Load reduces the number of steps to manually load data, see [Load Data Manually](#). (WL #15811)
- MySQL HeatWave now supports the following string functions:
  - `BIN( )`
  - `BIT_LENGTH( )`
  - `ELT( )`
  - `EXPORT_SET( )`
  - `FIELD( )`
  - `MAKE_SET( )`

These functions do not support [ENUM](#) type columns, see [Data Type Limitations](#). (WL #15628)

- MySQL HeatWave now supports full-table aggregation and group-by support for the following temporal data types:
  - `DATE`
  - `TIME`
  - `DATETIME`
  - `TIMESTAMP`
  - `YEAR`

MySQL HeatWave now supports the following functions with temporal data types:

- `SUM( )`
- `AVG( )`

- [STD\(\)](#)
- [STDDEV\(\)](#)
- [STDDEV\\_POP\(\)](#)
- [STDDEV\\_SAMP\(\)](#)
- [VAR\\_POP\(\)](#)
- [VAR\\_SAMP\(\)](#)
- [VARIANCE\(\)](#)

See: [Aggregate Functions](#). (WL #15637)

- MySQL HeatWave now enables the internal conversion of values from one data type to another, including all combinations of these different encodings:
  - Numeric values [INTEGER](#), [FLOAT](#), and [DECIMAL](#).
  - String values, excluding dictionary-encoded strings.
  - Temporal values [DATE](#), [TIME](#), [DATETIME](#), and [YEAR](#) data types.

This now permits MySQL HeatWave to offload queries that contain functions with unexpected data types by adding an internal cast from the unexpected data type to an expected data type. (WL #15465)

## Changes in MySQL HeatWave 8.1.0-u4 (2023-09-26, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- The MySQL HeatWave AutoML recommendation task now supports implicit feedback. See [Generate Recommendations](#). (WL #15802)

### MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse now uses schema inference to create an external table, see [Load Data from Object Storage](#). (WL #15851)
- MySQL HeatWave Lakehouse now includes support for MySQL HeatWave AutoML. See: [Use MySQL HeatWave AutoML with Lakehouse](#). (WL #15539)

### Functionality Added or Changed

- MySQL InnoDB on MySQL HeatWave now has support for a bulk load extension to the [LOAD DATA](#) statement. This is only available on MySQL HeatWave on AWS. See: [Bulk Ingest Data](#). (WL #14717, WL #14772, WL #15131, WL #15132, WL #15133, WL #15612)

## Changes in MySQL HeatWave 8.1.0-u3 (2023-09-12, General Availability)

- [MySQL HeatWave AutoML](#)
- [MySQL HeatWave Lakehouse](#)
- [Functionality Added or Changed](#)

## MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports text data types, see: [Supported Data Types for MySQL HeatWave AutoML](#). (WL #15743)

## MySQL HeatWave Lakehouse

- MySQL HeatWave Lakehouse can now use a regular expression to define a set of files. (WL #15840)
- MySQL HeatWave Lakehouse now uses Auto Parallel Load to load Parquet files. (WL #15656)
- MySQL HeatWave Lakehouse now supports the Avro file format, see [Supported File Formats and Data Types](#). (WL #15494)

## Functionality Added or Changed

- Adaptive Query Execution dynamically adjusts the query execution plan based on runtime statistics to improve query execution time. See: [About MySQL HeatWave](#). (WL #15368)

## Changes in MySQL HeatWave 8.1.0-u2 (2023-08-08, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

## MySQL HeatWave AutoML

- MySQL HeatWave AutoML progress tracking can now track the progress with finer granularity, and show the overall progress as well as the progress of subcomponents, see: [Track Progress for MySQL HeatWave AutoML Routines](#). (WL #15531)

## Functionality Added or Changed

- MySQL HeatWave improves time zone support. When the time zone is set to `SYSTEM` or set to an offset from UTC, queries run in MySQL HeatWave. `DATE`, `DATETIME`, `TIMESTAMP`, and `TIME` are processed based on the time-zone value. (Bug #35710134)

## Changes in MySQL HeatWave 8.1.0-u1 (2023-07-25, General Availability)

## Functionality Added or Changed

- MySQL HeatWave now supports the following functions:
  - `MID`
  - `REGEXP_INSTR`
  - `SOUNDEX`
  - `SPACE`
  - `WEIGHT_STRING`

(WL #15627)

## Changes in MySQL HeatWave 8.1.0 (2023-07-18, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML has enhanced recommendation systems that can now recommend new users who are similar to a specific user and new items that are similar to other items. See [Generate Recommendations](#). (WL #15674)
- Improvements to MySQL HeatWave AutoML models, the model catalog, and `ML_MODEL_IMPORT` now fully support ONNX and MySQL HeatWave AutoML model import. (WL #15383, WL #15559)

### Functionality Added or Changed

- Memory-aware query optimization now automatically adjusts MySQL HeatWave to optimize query execution time or memory usage and reduce the likelihood of queries running out of memory without user intervention. (WL #15529)
- It is now possible to track the memory used by change propagation and the number of transactions committed to MySQL InnoDB that are waiting to be propagated to MySQL HeatWave. (WL #15326, WL #15688)
- MySQL HeatWave now accelerates queries with `INTERSECT` and `EXCEPT` clauses. With this support, MySQL HeatWave can now offload and accelerate all 97 TPC-DS queries. (WL #15362)
- MySQL HeatWave now supports the following temporal functions:
  - `GET_FORMAT()`
  - `MAKETIME()`
  - `PERIOD_ADD()`
  - `PERIOD_DIFF()`
  - `SEC_TO_TIME()`
  - `SUBTIME()`
  - `SYSDATE()`
  - `TIMEDIFF()`

(WL #14958)

## Changes in MySQL HeatWave 8.0.33-u3 (2023-06-09, General Availability)

### MySQL HeatWave Lakehouse

- MySQL HeatWave expands to include MySQL HeatWave Lakehouse, letting organizations process and query hundreds of terabytes of data residing in Object Storage—in a variety of file formats, such as CSV and Parquet.

The Lakehouse feature of MySQL HeatWave enables query processing on Object Storage-resident data. The source data is read from Object Storage, transformed to the MySQL HeatWave format, stored in the MySQL HeatWave persistence storage layer in OCI Object Storage, and then loaded to MySQL HeatWave Cluster memory.

- Provides in-memory query processing on Object Storage-resident data.
- Data is not loaded into the MySQL InnoDB storage layer.
- Supports structured and relational data in CSV and Parquet formats.
- With this feature, users can now analyze data in both MySQL InnoDB and an object store using familiar SQL syntax in the same query.

See [Load Data from Object Storage into MySQL HeatWave Cluster](#). (WL #15575)

## Changes in MySQL HeatWave 8.0.33 (2023-04-27, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports a recommendation task, see [Generate Recommendations](#). (WL #15416)

### Functionality Added or Changed

- The new MySQL HeatWave Auto Unload utility automates the process of unloading tables from MySQL HeatWave. (WL #15447)
- The following columns were added to the `performance_schema.rpd_query_stats` table:
  - `CONNECTION_ID`: The ID of the connection.
  - `STATEMENT_ID`: The global query ID.

(WL #15526)

## Changes in MySQL HeatWave 8.0.32-u2 (2023-02-21, General Availability)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML now supports unsupervised anomaly detection, which is the data mining task that finds unusual patterns in data. (WL #15518)

## Changes in MySQL HeatWave 8.0.32-u1 (2023-02-21, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML adds support for multivariate endogenous forecasting models, and exogenous forecasting models. (WL #15511)

## Functionality Added or Changed

- The `rpd_tables` table now supports recovery time measurement with `RECOVERY_SOURCE`, `RECOVERY_START_TIMESTAMP` and `RECOVERY_END_TIMESTAMP`. The `performance_schema.global_status` now includes the `rapid_recovery_time` variable. (WL #15441)
- Arithmetic operators are now supported with variable-length encoded string columns. Mathematical functions are now supported with variable-length columns. (WL #15405)

## Changes in MySQL HeatWave 8.0.32 (2023-01-17, General Availability)

- [Deprecation and Removal Notes](#)
- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

## Deprecation and Removal Notes

- MySQL 8.0.32 deprecates the `heatwave_load_report` and `heatwave_advisor_report` tables, and replaces them with the `heatwave_autopilot_report` table in the `sys` schema. A future release will remove them. (Bug #34727481)

## MySQL HeatWave AutoML

- MySQL 8.0.32 introduces a number of changes for MySQL HeatWave AutoML routines:
    - The routines: `ML_PREDICT_ROW`, `ML_PREDICT_TABLE`, `ML_EXPLAIN_ROW`, and `ML_EXPLAIN_TABLE` now include the `ml_results` column, which uses `JSON` format to return the results.
    - The routines: `ML_PREDICT_ROW`, `ML_PREDICT_TABLE`, `ML_EXPLAIN_ROW`, and `ML_EXPLAIN_TABLE` now allow additional columns that are not required for prediction or explanation.
    - The routines: `ML_PREDICT_ROW`, and `ML_PREDICT_TABLE` now allow an `options` parameter in `JSON` format.
    - The `ML_TRAIN` routine also runs the `ML_EXPLAIN` routine with the default Permutation Importance model.
- (WL #15420)
- MySQL 8.0.32 introduces progress tracking for `ML_TRAIN`. Use a second MySQL Client window to track the progress of `ML_TRAIN` with calls to the performance schema. It also introduces the `rapid_ml_operation_count` status variable. (WL #15384)

## Functionality Added or Changed

- MySQL HeatWave now supports the following encryption and compression functions:
  - `COMPRESS()`
  - `MD5()`
  - `RANDOM_BYTES()`

- [SHA\(\)](#)
- [SHA1\(\)](#)
- [SHA2\(\)](#)
- [UNCOMPRESS\(\)](#)
- [UNCOMPRESSED\\_LENGTH\(\)](#)

(WL #15283)

- MySQL HeatWave now supports the following mathematical functions:

- [CRC32\(\)](#)
- [LOG2\(\)](#)
- [LOG10\(\)](#)
- [RAND\(\)](#)

(WL #15256)

## Changes in MySQL HeatWave 8.0.31 (2022-10-11, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave AutoML queries are now monitored and recorded in the Performance Schema tables [rpd\\_query\\_stats](#) and [rpd\\_exec\\_stats](#). Where a single MySQL HeatWave AutoML query contains a number of sub-queries, there is one record in [rpd\\_query\\_stats](#) and multiple records in [rpd\\_exec\\_stats](#). (WL #15243)
- New functions have been added to MySQL HeatWave AutoML to help you manage models:
  - When you run the [ML\\_TRAIN](#) routine on a training dataset, you can now specify a model handle to use for the model instead of the generated one.
  - A new column [notes](#) has been added to the [MODEL\\_CATALOG](#) table, which you can use to record notes about the models in your model catalog.
  - The new column [model\\_metadata](#) in the [MODEL\\_CATALOG](#) table records metadata for models, such as the training score, training time, and information about the training dataset. If an error occurs during training or you cancel the training operation, MySQL HeatWave AutoML records the error status in this column.

(WL #15243)

- MySQL HeatWave AutoML now supports the upload of pre-trained models in ONNX (Open Neural Network Exchange) format to the model catalog. You can load them using the stored procedure [ML\\_MODEL\\_IMPORT](#) that provides the conversion required to store the model in a MySQL table. (WL #15243)

- A new stored procedure [ML\\_EXPLAIN](#) lets you train a variety of model explainers and prediction explainers for MySQL HeatWave AutoML, in addition to the default Permutation Importance model and prediction explainers:
  - The Partial Dependence model explainer shows how changing the values of one or more columns will change the value that the model predicts.
  - The SHAP model explainer produces global feature importance values based on Shapley values.
  - The Fast SHAP model explainer is a subsampling version of the SHAP model explainer which usually has a faster runtime.
  - The Permutation Importance prediction explainer explains the prediction for a single row or table.
  - The SHAP prediction explainer uses feature importance values to explain the prediction for a single row or table.

When you use the [ML\\_EXPLAIN\\_TABLE](#) and [ML\\_EXPLAIN\\_ROW](#) stored procedures to generate explanations for a prediction, you can now use the SHAP prediction explainer as an alternative to the default Permutation Importance prediction explainer. SHAP produces feature importance values (explanations) based on Shapley values. (WL #15243)

- MySQL HeatWave AutoML now supports timeseries forecasting using the existing stored procedures [ML\\_TRAIN](#), [ML\\_PREDICT\\_TABLE](#), and [ML\\_SCORE](#). You can create a forecast for a single column (a univariate endogenous variable) with a numeric data type. The forecasting task is specified as a JSON object when you call the [ML\\_TRAIN](#) stored procedure. (WL #15243)

## Functionality Added or Changed

- MySQL HeatWave uses dictionary encoding to compress string columns (CHAR, VARCHAR, TEXT). These dictionaries are built for each string column with the [RAPID\\_COLUMN=ENCODING=SORTED](#) keyword. MySQL HeatWave now supports 8.5 billion dictionary entries (up from 4 billion), which means MySQL HeatWave can now encode string columns with number of distinct value (NDV) up to 8.5 billion. (WL #14742)
- MySQL HeatWave now uses statistics based on minimum and maximum values to create zone maps for every primary key column. MySQL HeatWave then uses the zone maps for range and point queries to only scan data chunks that are relevant for the query, and accelerates these queries by an order of magnitude. This is particularly useful for improving range queries in OLAP and mixed workloads. (WL #14713)
- A new MySQL optimizer is introduced for MySQL HeatWave to provide a holistic cost model across MySQL and MySQL HeatWave, create better query plans based on statistics used in MySQL Autopilot, reduce compilation time, eliminate the need of query hints for join order, and improve join query performance. With the new optimizer, MySQL HeatWave can now run all 22 TPC-H queries without straight join hints. Before 8.0.31, a straight join hint is needed for 10 out of 22 TPC-H to reach peak performance. (WL #14449)
- DDL statements such as [ALTER TABLE](#), [RENAME TABLE](#), and [TRUNCATE TABLE](#) are now permitted on a table that has RAPID defined as the secondary engine. If a DDL operation is successfully carried out on a table that is loaded to a MySQL HeatWave Cluster at the time, MySQL HeatWave automatically reloads the table from MySQL InnoDB. Note that if the DDL operation makes the table's structure incompatible with MySQL HeatWave, the table is unloaded from the MySQL HeatWave Cluster. (WL #15129)

## Changes in MySQL HeatWave 8.0.30-u1 (2022-09-06, General Availability)

## Functionality Added or Changed

- MySQL HeatWave now supports the [ABS\(\)](#), [POWER\(\)](#), and [SIGN\(\)](#) mathematical functions. (WL #15246)
- MySQL HeatWave now supports the [GROUP\\_CONCAT\(\)](#) aggregation function with variable-length ([VARLEN](#)) string columns. (WL #14767)
- Change propagation memory management was enhanced to improve MySQL HeatWave shutdown time. (WL #15306)
- Query performance was optimized for compressed data. Decompression of data that is queried but does not participate in filter evaluation is delayed until after a filter is applied. With this optimization, decompression is avoided entirely for data that is filtered out. (WL #15169)

## Changes in MySQL HeatWave 8.0.30 (2022-07-26, General Availability)

- [Advisor](#)
- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### Advisor

- MySQL HeatWave Advisor Auto Encoding, which recommends string column encodings, now provides encoding recommendations that optimize query performance. Recommendations are based on performance models that use query execution data. Previously, string column encoding recommendations were optimized for cluster memory usage only. A performance improvement estimate is provided with string column encoding recommendations. (Bug #34145862)

### MySQL HeatWave AutoML

- You can now train MySQL HeatWave AutoML models on tables containing [DATE](#), [TIME](#), [DATETIME](#), [TIMESTAMP](#), and [YEAR](#) data types. (Bug #33895503)
- MySQL HeatWave AutoML now generates a model explanation when you train a machine learning model. Model explanations help identify the features that are most important to a model. For more information, see [The Model Catalog](#).

The following columns were added to the [MODEL\\_CATALOG](#) table:

- [column\\_names](#): The feature columns used to train the model.
- [last\\_accessed](#): The last time the model was accessed. MySQL HeatWave AutoML routines update this value to the current timestamp when accessing the model.
- [model\\_explanation](#): The model explanation generated during training.
- [model\\_type](#): The type of model (algorithm) selected by [ML\\_TRAIN](#) to build the model.
- [task](#): The task type specified in the [ML\\_TRAIN](#) query ([classification](#) or [regression](#)).

[ML\\_PREDICT \\*](#) and [ML\\_EXPLAIN \\*](#) routine performance was improved, resulting in faster prediction and explanation processing. (WL #15088, WL #15014)

- The following MySQL HeatWave AutoML enhancements were implemented:
  - `ML_TRAIN` options for advanced users. These options permit users to customize various aspects of the ML training pipeline including algorithm selection, feature selection, and hyperparameter optimization.
  - The `model_list` option permits specifying the type of model to be trained.
  - The `exclude_model_list` option specifies models types to exclude from consideration during model selection.
  - The `optimization_metric` option specifies the scoring metric to optimize for when training a machine learning model.
  - The `exclude_column_list` option specifies feature columns to exclude from consideration when training a machine learning model.

For more information, see [ML\\_TRAIN](#).

- Support was added for Support Vector Machine `SVC` and `LinearSVC` classification and regression models. For a complete list of supported model types, see [Model Types](#).
- The `ML_TRAIN` routine now reports a message if a trained model does not meet expected quality criteria.
- `ML_EXPLAIN_ROW` and `ML_EXPLAIN_TABLE` routines now provide information to help interpret explanations. The routines also report a warning when a model quality issue is detected, enabling users to revisit their data in order to improve model quality.

(WL #15089)

## Functionality Added or Changed

- The amount of heap memory allocated on the MySQL node for each table loaded into MySQL HeatWave was reduced, increasing the maximum number of tables that can be loaded. For `MySQL.HeatWave.VM.E3.Standard` shapes, the maximum was raised from 100k tables to 400k tables. For `MySQL.HeatWave.BM.E3.Standard` shapes, the maximum number was raised from 400k tables to 1600k tables. The actual number of tables that can be loaded is dependent on the table's data. (Bug #33951708)
- The `performance_schema.rpd_column_id` table was modified to remove redundant data. The `NAME`, `SCHEMA_NAME`, `TABLE_NAME` columns were removed, and a `TABLE_ID` column was added. (Bug #33899183)
- Support was added for the `FROM_DAYS( )` temporal function, and `GREATEST( )` and `LEAST( )` comparison and string functions which now support `DATE`, `DATETIME`, `TIME`, and `TIMESTAMP` columns. (WL #14956)
- Support was added for built-in server-side data masking and de-identification to help protect sensitive data from unauthorized uses by hiding and replacing real values with substitutes. Data masking and de-identification operations are performed on the server, and queries involving data masking and de-identification functions are accelerated by MySQL HeatWave. The following data masking and de-identification functions are supported:
  - `gen_range( )`
  - `gen_rnd_email( )`

- `gen_rnd_pan()`
- `gen_rnd_ssn()`
- `gen_rnd_us_phone()`
- `mask_inner()`
- `mask_outer()`
- `mask_pan()`
- `mask_pan_relaxed()`
- `mask_ssn()`

See [Data Masking and De-Identification Functions](#). (WL #15143)

- Optimizations were implemented to improve performance for `JOIN` and `GROUP BY` queries with execution plans involving multiple consecutive rounds of data partitioning. (WL #15143)
- `expr IN (value,...)` comparisons, where the expression is a single value and compared values are constants of the same data type and encoding, have been optimized. For example, the following `IN()` comparison has been optimized:

```
SELECT * FROM Customers WHERE Country IN ('Germany', 'France', 'Spain');
```

(WL #14952)

## Changes in MySQL HeatWave 8.0.28-u3 (2022-04-19, General Availability)

### Functionality Added or Changed

- Tables that have become stale due to a change propagation failure resulting from an out-of-code error are now automatically reloaded. A check for stale tables is performed periodically when the MySQL HeatWave Cluster is idle. Previously, identifying and reloading stale tables was a manual process. See [About Change Propagation](#). (WL #14914)

## Changes in MySQL HeatWave 8.0.28-u2 (2022-03-29, General Availability)

- [MySQL HeatWave AutoML](#)
- [Functionality Added or Changed](#)

### MySQL HeatWave AutoML

- MySQL HeatWave customers now have access to MySQL HeatWave AutoML, which is a fully managed, highly scalable, cost-efficient, machine learning solution for data stored in MySQL. MySQL HeatWave AutoML provides a simple SQL interface for training and using predictive machine learning models, which can be used by novice and experienced ML practitioners alike. With MySQL HeatWave AutoML, you can train a model with a single call to an SQL routine. Similarly, you can generate predictions with a single `CALL` or `SELECT` statement which can be easily integrated with your applications.

With MySQL HeatWave AutoML, data and models never leave the MySQL Database Service, saving you time and effort while keeping your data and models secure. MySQL HeatWave AutoML is optimized for MySQL HeatWave shapes and scaling, and all MySQL HeatWave AutoML processing is performed on the MySQL HeatWave Cluster. ML computation is distributed among MySQL HeatWave nodes,

taking advantage of MySQL HeatWave's scalability and massively parallel processing capabilities. For more information about MySQL HeatWave's machine learning capabilities, see [Train and Use Machine Learning Models](#). (WL #14661, WL #14836, WL #15014)

## Functionality Added or Changed

- MySQL HeatWave now compresses data as it is loaded, which permits MySQL HeatWave nodes to store more data. More data per node reduces costs by minimizing the size of the MySQL HeatWave Cluster required to store your data. Data compression is enabled by default but can be disabled at runtime using the [rapid\\_compression](#) session variable. For more information, see [Enable or Disable Data Compression](#). (WL #14868)

## Changes in MySQL HeatWave 8.0.28-u1 (2022-02-15, General Availability)

### Functionality Added or Changed

- MySQL HeatWave now supports up to 1017 columns for base relations (tables as loaded into MySQL HeatWave), and up to 1800 columns for intermediate relations (intermediate tables used during query processing). The maximum column width for base relations and intermediate relations was increased to 65532 bytes. See [Column Limits](#). (WL #14918)

## Changes in MySQL HeatWave 8.0.27-u3 (2021-12-15, General Availability)

### Functionality Added or Changed

- Support was added for the [CONVERT\\_TZ\( \)](#) and [LAST\\_DAY\( \)](#) functions, which are used to manipulate temporal values. (WL #14768)
- The MySQL HeatWave Cluster recovery process was optimized to avoid applying a large volume of changelogs during recovery. Snapshots are now taken when the volume of changelogs and the time required to apply those changes exceed specific thresholds, and recovery from Object Storage is performed using those snapshots. (WL #14615)
- MySQL HeatWave now supports automatic data reload when the MySQL HeatWave Cluster is restarted. Previously, when a MySQL HeatWave Cluster was stopped by a stop or restart action, data had to be reloaded manually after the cluster was restarted. Now, when starting or restarting a MySQL HeatWave Cluster, data that was previously loaded is reloaded automatically. Data changes that occur on the DB System while the MySQL HeatWave Cluster is offline are included in the reloaded data.

Automatic data reload does not occur if the MySQL HeatWave Cluster was stopped as a result of a stop or restart action performed on the DB System to which the MySQL HeatWave Cluster is attached. In this case, data loaded in the MySQL HeatWave Cluster must be reloaded manually after the MySQL HeatWave Cluster is restarted. (WL #14729)

## Changes in MySQL HeatWave 8.0.27-u2 (2021-12-07, General Availability)

### Functionality Added or Changed

- MySQL HeatWave now supports querying views. See [Query Views](#). (WL #13568)
- Bloom filter join optimizations were introduced. For MySQL HeatWave queries that join large and small relations, bloom filters reduce the amount of data processed by early filtering and the amount of memory used during query processing. (WL #14752)
- The [rpd\\_query\\_stats](#) table, which stores MySQL HeatWave query history (compilation and execution statistics), now stores data for the last 1000 executed queries. Previously, data was stored for the last 200 queries.

The following columns were added to the `performance_schema.rpd_tables` table:

- `SIZE_BYTES`: The amount of data loaded per table, in bytes.
- `QUERY_COUNT`: The number of queries that referenced the table.
- `LAST_QUERIED`: The timestamp of the last query that referenced the table.
- `LOAD_END_TIMESTAMP`: The load completion timestamp for the table.

The following column was added to the `performance_schema.rpd_columns` table:

- `DICT_SIZE_BYTES`: The dictionary size per column, in bytes.

The following column was added to the `performance_schema.rpd_nodes` table:

- `BASEREL_MEMORY_USAGE`: The base relation memory footprint per node.

The `rapid_query_stats` and `rpd_exec_stats` tables are now synchronized. If a query record is removed from the `rapid_query_stats` table, it is also removed from the `rpd_exec_stats` table. (WL #14759)

## Changes in MySQL HeatWave 8.0.26-u2 (2021-09-21, General Availability)

### Functionality Added or Changed

- The following function support was added:
  - `YEARWEEK(date)`, `YEARWEEK(date,mode)`
  - The mode argument for the two-argument form of the `WEEK()` function: `WEEK(date[,mode])`
  - `MAKEDATE()`
  - “Zero” handling for dates such as '2001-11-00' was implemented for `WEEK()`, `YEARWEEK()`, and `MAKEDATE()` functions.
  - `CAST()` of `FLOAT` and `DOUBLE` values to `DECIMAL`

(Bug #33163625, Bug #33138534, WL #14714)

- The new `hw_data_scanned` global status variable tracks the total cumulative megabytes scanned by successfully executed MySQL HeatWave queries.

The number of megabytes scanned by an individual MySQL HeatWave query can be obtained by querying the `performance_schema.rpd_query_stats` table.

An estimated number of megabytes scanned by an individual query can be obtained by running the query with `EXPLAIN` and querying the `performance_schema.rpd_query_stats` table.

For more information, see [Scanned Data Monitoring](#). (WL #14738)

## Changes in MySQL HeatWave 8.0.26-u1 (2021-08-10, General Availability)

- [MySQL HeatWave Network Layer](#)
- [MySQL HeatWave Data Management Layer](#)

- [Functionality Added or Changed](#)

## MySQL HeatWave Network Layer

- MySQL HeatWave network layer optimizations have improved scalability and network performance. (WL #14513)

## MySQL HeatWave Data Management Layer

- Data loaded into MySQL HeatWave, including propagated changes, are now persisted to OCI Object Storage for recovery in case of a MySQL HeatWave node or cluster failure. Previously, data was recovered from the MySQL DB System. Loading data from OCI Object Storage is faster because data does not need to be converted to the MySQL HeatWave storage format, as is required when loading data from the MySQL DB System. If data recovery from OCI Object Storage fails, MySQL HeatWave falls back to recovering data from the MySQL DB System. Data removed from MySQL HeatWave when a table is unloaded is removed from OCI Object Storage in a background operation. For related information, see [MySQL HeatWave Cluster Failure and Recovery](#). (WL #14478, WL #14046, WL #14541)

## Functionality Added or Changed

- MySQL HeatWave now supports `COUNT(NULL)`, except in cases where it is used as an input argument for non-aggregate operators. (Bug #33005146)
- Full support was added for the `DISTINCT` modifier. Previously, multiple instances of (`DISTINCT value`) expressions in a query were only permitted if the same `value` was specified. (Bug #32865043, Bug #33007714, WL #14574)
- MySQL HeatWave now supports the `WITH ROLLUP` modifier in `GROUP BY` clauses. (WL #14533)
- MySQL HeatWave now supports window functions. For optimal performance, window functions in MySQL HeatWave utilize a massively parallel, partitioning-based algorithm. For more information, see [Window Functions](#). (WL #14674)

## Changes in MySQL HeatWave 8.0.26 (2021-07-23, General Availability)

- [Advisor](#)
- [Auto Parallel Load](#)
- [Auto Scheduling](#)
- [Functionality Added or Changed](#)

### Advisor

- The new MySQL HeatWave Advisor provides string column encoding and data placement key recommendations based on machine learning models, data analysis, and MySQL HeatWave query history. Implementing MySQL HeatWave Advisor recommendations can improve query performance and reduce the amount of memory required on MySQL HeatWave nodes.

The MySQL HeatWave Advisor also provides a Query Insights feature, which provides runtimes for successfully executed queries, and runtime estimates for `EXPLAIN` queries, queries canceled using `Ctrl+C`, and queries that fail due to out of memory errors. Runtime data is useful for query optimization, troubleshooting, and estimating the cost of running a particular query or workload.

The MySQL HeatWave Advisor is implemented as a stored procedure named `heatwave_advisor`, which resides in the MySQL `sys` schema. Running Advisor involves issuing a `CALL` statement for the stored procedure with optional arguments.

```
CALL sys.heatwave_advisor (options);
```

For more information about the MySQL HeatWave Advisor, see [Optimize Workloads for OLAP](#). (WL #14328, WL #14510, WL #14431, WL #14328, WL #14651)

## Auto Parallel Load

- The new MySQL HeatWave Auto Parallel Load utility automates the process of preparing and loading tables into MySQL HeatWave and loads data using an optimized number of parallel load threads.

The MySQL HeatWave Auto Parallel Load utility is implemented as a stored procedure named `heatwave_load`, which resides in the MySQL `sys` schema. Running Auto Parallel Load involves issuing a `CALL` statement for the stored procedure, which takes a list of schemas and options as arguments.

```
CALL sys.heatwave_load (db_list, [options]);
```

For more information about the MySQL HeatWave Auto Parallel Load utility, see [Load Data Using Auto Parallel Load](#). (WL #14149)

## Auto Scheduling

- The MySQL HeatWave query scheduling algorithm was improved. The revised algorithm prioritizes queries based on estimated cost and wait time in the queue, which enables dynamic, workload-aware query prioritization. Previously, queries were prioritized using a static cost-based prioritization model. (WL #14608)

## Functionality Added or Changed

- `DATE_ADD( )` and `DATE_SUB( )` functions now support precision `INTERVAL` values (`DECIMAL`, `DOUBLE`, and `FLOAT`). (Bug #32725985, Bug #32438123)
- Support was added for multiple instances of `COUNT(DISTINCT)` in a query. (Bug #32422984)
- Query compilation and processing was improved to permit combining aggregate operators into a single task in the physical query plan, which avoids fully materializing intermediate result sets. This enhancement reduces memory allocation and deallocation operations, memory usage, and execution time for affected queries. (WL #14614)
- The cost model that estimates MySQL HeatWave query runtimes can now use statistics from previously executed queries, which improves the accuracy of query runtime estimates. (WL #14546)
- MySQL HeatWave now supports `CREATE TABLE ... SELECT` statements where the `SELECT` query is offloaded to MySQL HeatWave and the table is created on the MySQL Database Service instance. This feature improves `CREATE TABLE ... SELECT` performance in cases where the `SELECT` portion of the statement is a long running, complex query. For more information, see [CREATE TABLE ... SELECT Statements](#). (WL #14516)
- Support was added for `REGEXP_REPLACE( )` and `REGEXP_SUBSTR( )` regular expression functions, and error messaging was improved for `REGEXP( )` function syntax mismatches, expression errors, and input argument errors. (WL #14641)

## Changes in MySQL HeatWave 8.0.25 (2021-05-11, General Availability)

## Functionality Added or Changed

- Support was added for `CAST( )` of `ENUM` column values to `CHAR` or `VARCHAR` where the `ENUM` value is cast to a `FLOAT` value, as in the following example:

```
SELECT CAST(CAST(enum_col AS FLOAT) AS CHAR(3)) FROM tbl_name;
```

(Bug #32618454)

- Support was added for `SELECT DISTINCT` queries that order the result set by a column that is not defined in the `SELECT` list. For example, the following query can now be offloaded to MySQL HeatWave for execution:

```
SELECT DISTINCT a FROM t1 ORDER BY c DESC;
```

(Bug #32583856)

- Query plan statistics are now collected and stored in a statistics cache when a query is executed in MySQL HeatWave. When a new query shares query execution plan nodes with previously executed queries, the actual statistics collected from previously executed queries are used instead of estimated statistics, which improves query execution plans, cost estimations, execution times, and memory efficiency.

The statistics cache is an LRU structure. When cache capacity is reached, the least recently used entries are evicted from the cache as new entries are added. The maximum number of entries permitted in the statistics cache is defined by the `rapid_stats_cache_max_entries` setting. The number of entries permitted by default is 65536, which is enough to store statistics for 4000 to 5000 unique queries of medium complexity. (WL #14503)

- Support was added for:
  - `CAST( ) AS YEAR`. Both variable-length and dictionary-encoded string columns values are supported.
  - The `FORMAT( )` function. Variable-length-encoded string columns are not supported.

(WL #14511)

## Changes in MySQL HeatWave 8.0.24 (2021-04-20, General Availability)

### Functionality Added or Changed

- Comparison of different temporal type values is now supported. For example, a query that compares `DATE` values to `TIMESTAMP` values can now be offloaded to MySQL HeatWave. (Bug #32420986)
- Range operators on `VARLEN`-encoded string columns are now supported. For example, the following query, where `L_LINESTATUS` is a `VARLEN`-encoded string column, can now be offloaded to MySQL HeatWave:

```
SELECT COUNT(*) FROM lineitem WHERE L_LINESTATUS >= 1 and L_LINESTATUS <= 10;
```

(Bug #31721399)

- MySQL HeatWave now supports `INSERT ... SELECT` statements where the `SELECT` query is offloaded to MySQL HeatWave and the result set is inserted into a table on the MySQL Database Service instance. This feature improves `INSERT ... SELECT` performance in cases where the `SELECT` portion of the statement is a long running, complex query. For more information, see [INSERT ... SELECT Statements](#). (WL #14299)

- [VARLEN](#)-encoded columns are now supported as data placement keys. For information about the data placement feature, see [Data Placement Keys](#). (WL #14491)
- Failure handling was improved for queries involving unsupported internal data types. Such queries now exit with an error indicating that the internal data type of the query is not supported. (WL #14483)

## Changes in MySQL HeatWave 8.0.23-u2 (2021-03-15, General Availability)

### Functionality Added or Changed

- Support was added for the following aggregate functions:

- [STD\(\)](#)
- [STDDEV\(\)](#)
- [STDDEV\\_POP\(\)](#)
- [STDDEV\\_SAMP\(\)](#)
- [VAR\\_POP\(\)](#)
- [VAR\\_SAMP\(\)](#)
- [VARIANCE\(\)](#)

See [Aggregate Functions](#). (WL #14479)

- MySQL HeatWave now uses a priority-based scheduling mechanism based on query cost estimates to schedule queries for execution. Previously, queries were executed in the order of arrival. The scheduling mechanism prioritizes short running queries over long running queries to reduce overall query execution wait times. For more information, see [Auto Scheduling](#). (WL #14423)

## Changes in MySQL HeatWave 8.0.23-u1 (2021-02-09, General Availability)

### Functionality Added or Changed

- String column encoding support was added for [TEXT](#)-type columns. See [Encoding String Columns](#). (WL #14430)
- [UNION](#) and [UNION ALL](#) support was extended. The clauses are now supported at any location in a query that is permitted by MySQL. (WL #14455)
- The following temporal functions are now supported:
  - [TO\\_SECONDS\(\)](#)
  - [UNIX\\_TIMESTAMP\(\)](#)
  - [FROM\\_UNIXTIME\(\)](#)
  - [TIME\\_TO\\_SEC\(\)](#)

See [Temporal Functions](#).

The following temporal functions are now supported with [VARLEN](#)-encoded columns:

- [TO\\_DAYS\(\)](#)

- `DAYOFYEAR()`
- `QUARTER()`
- `TO_SECONDS()`

See [Temporal Functions](#).

The following string functions are now supported with `VARLEN`-encoded columns:

- `ORD()`
- `ASCII()`

See [String Functions and Operators](#).

`SET timezone = timezone` with the `timezone` value specified as an offset from UTC in the form of `[H]H:MM` and prefixed with a `+` or `-` is now supported with the `UNIX_TIMESTAMP()` and `FROM_UNIXTIME()` functions. (WL #14345)

- Offset is now supported with the `LIMIT` clause:

```
SELECT * FROM tbl LIMIT offset,row_count;
```

The PostgreSQL syntax is also supported:

```
SELECT * FROM tbl LIMIT row_count OFFSET offset;
```

(WL #14341)

- New Performance Schema tables provide access to query and execution statistics:

- `performance_schema.rpd_exec_stats`
- `performance_schema.rpd_query_stats`

Changes to MySQL HeatWave Performance Schema tables:

- An `NDV` (Number of Distinct Values) column was added to the `performance_schema.rpd_columns` table.
- A `ROWS` column that shows the total number of rows in a table was added to the `performance_schema.rpd_tables` table.
- A `MEMORY_USAGE` column that shows node memory usage was added to the `performance_schema.rpd_columns` table.
- The `performance_schema.rpd_nodes` `DRAM` column was renamed to `MEMORY_TOTAL`. The `MEMORY_TOTAL` column shows the total memory allocated to a MySQL HeatWave node.

See [MySQL HeatWave Performance Schema Tables](#). (WL #14386)

## Index

### A

`ABS()`, 30

Adaptive Query Execution, 24  
Advisor, 31, 36  
Aggregate functions, 20, 22, 39  
ALTER TABLE, 29  
ASCII(), 39  
Auto Data Compression, 20, 22  
Auto Encoding, 31  
Auto Parallel Load, 14, 15, 20, 21, 22, 24, 24, 36  
Auto Unload, 20, 27  
AutoML, 4, 6, 10, 12, 13, 15, 17, 20, 21, 24, 24, 25, 26, 27, 27,  
27, 28, 29, 31, 33  
Autopilot, 14, 17  
AVG(), 23  
Avro file format, 22, 24

## **B**

BIN(), 23  
BIT\_LENGTH(), 23  
Bloom filter optimization, 34  
Bulk load, 17, 20, 24

## **C**

CAST(), 23, 35, 37  
Change propagation, 26, 31, 33  
Column limits, 34  
Comparison functions , 21, 31  
Comparisons, 31  
COMPRESS(), 28  
Compression functions, 28  
CONVERT\_TZ(), 34  
COUNT(DISTINCT), 36  
COUNT(NULL), 35  
CRC32(), 28  
CREATE TABLE ... SELECT, 36  
CSV file format, 26  
CUBE, 22

## **D**

Data compression, 31, 33  
Data placement, 38  
Data types, 38  
DATE, 23, 25  
DATETIME, 23, 25  
DATE\_ADD(), 36  
DATE\_SUB(), 36  
DAYOFYEAR(), 39  
Dictionary encoding, 29  
DISTINCT, 35

## **E**

ELT(), 23  
Encryption functions, 28  
Error reporting, 15, 21

EXCEPT, 26  
EXPORT\_SET(), 23

## F

FIELD(), 23  
FORMAT(), 37  
FROM\_DAYS(), 31  
FROM\_UNIXTIME(), 39

## G

GenAI, 4, 6, 7, 8, 10, 11, 12, 14, 15, 17  
gen\_range(), 31  
gen\_rnd\_email(), 31  
gen\_rnd\_pan(), 31  
gen\_rnd\_ssn(), 31  
gen\_rnd\_us\_phone(), 31  
GET\_FORMAT(), 26  
GREATEST(), 21, 31  
GROUP BY, 22, 35  
GROUP\_CONCAT(), 31  
Guided Load, 20, 23

## H

Heap segment size, 31  
heatwave\_advisor\_report table, 28  
heatwave\_autopilot\_report table, 28  
heatwave\_load\_report table, 28  
High Availability, 20  
HLL(), 20, 22  
HyperLogLog, 20, 22

## I

IN(), 31  
INSERT ... SELECT, 38  
INTERSECT, 26

## J

JavaScript, 14  
JSON ->, 22  
JSON ->>, 22  
JSON column path, 22  
JSON data type, 20  
JSON file format, 20  
JSON functions, 22  
JSON inline path, 22  
JSON\_ARRAY(), 22  
JSON\_DEPTH(), 22  
JSON\_EXTRACT(), 22  
JSON\_LENGTH(), 22  
JSON\_OBJECT(), 22  
JSON\_UNQUOTE(), 22

## L

Lakehouse, 4, 6, 7, 8, 10, 11, 12, 13, 14, 15, 17, 20, 20, 21, 21,  
22, 24, 24, 26  
Lakehouse Auto Parallel Load, 20, 21, 22, 24, 24  
Lakehouse Incremental Load, 17  
LAST\_DAY, 34  
LEAST(), 21, 31  
LIMIT, 39  
Loading data, 31  
LOG10(), 28  
LOG2(), 28

## M

Machine learning, 4, 6, 10, 12, 13, 15, 17, 20, 21, 24, 24, 25, 26, 27,  
27, 27, 28, 29, 31, 33  
MAKEDATE(), 35  
MAKETIME(), 26  
MAKE\_SET(), 23  
mask\_inner(), 31  
mask\_outer(), 31  
mask\_pan(), 31  
mask\_pan\_relaxed(), 31  
mask\_ssn(), 31  
Mathematical functions, 27, 28, 31  
MD5(), 28  
Memory usage, 26  
MID, 25  
MySQL, 4, 7, 8, 10, 10, 11, 14, 15, 17  
MySQL additional functionality, 17, 20, 20, 22, 24  
MySQL HeatWave, 4, 7, 8, 10, 10, 11, 14, 15, 17  
MySQL HeatWave AutoML, 4, 6, 10, 12, 13, 15, 17, 20, 21, 24, 24, 25, 26,  
27, 27, 27, 28, 29, 31, 33  
MySQL HeatWave Autopilot, 14, 17  
MySQL HeatWave GenAI, 4, 6, 7, 8, 10, 11, 12, 14, 15, 17  
MySQL HeatWave Lakehouse, 4, 6, 7, 8, 10, 11, 12, 13, 14, 15, 17, 20,  
20, 21, 21, 22, 24, 24, 26  
MySQL InnoDB, 17, 20, 24

## N

Networking, 35

## O

Object Storage, 34  
OFFSET, 39  
Optimizer, 29  
ORD(), 39  
Outbound Replication, 20

## P

Parquet file format, 24, 26  
Partitioning, 31  
Partitions, 14

Performance Schema, 27, 27, 31, 34, 39  
PERIOD\_ADD(), 26  
PERIOD\_DIFF(), 26  
Planned shutdown, 34  
Point-in-time-recovery, 20  
POWER(), 31

## Q

QUALIFY, 22  
QUARTER(), 39  
Queries, 35, 38  
Query cost model, 36  
Query optimization, 26  
Query plans, 29  
Query scheduling, 36, 39  
Query support, 21

## R

RAND(), 28  
RANDOM\_BYTES(), 28  
rapid\_stats\_cache\_max\_entries, 37  
Read Replication, 20  
Recovery, 34, 35  
REGEXP\_INSTR, 25  
REGEXP\_REPLACE(), 36  
REGEXP\_SUBSTR(), 36  
Regular expression functions, 36  
Reload tables, 20  
RENAME TABLE, 29  
rpd\_columns table, 39  
rpd\_column\_id table, 31  
rpd\_exec\_stats table, 39  
rpd\_nodes table, 39  
rpd\_query\_stats table, 27, 34, 39  
rpd\_tables table, 27

## S

Scalability, 35  
Schema inference, 20, 22, 24  
SEC\_TO\_TIME(), 26  
SELECT, 20, 22  
SELECT DISTINCT, 37  
SET timezone, 39  
SHA(), 28  
SHA1(), 28  
SHA2(), 28  
SIGN(), 31  
SOUNDEX, 25  
SPACE, 25  
Statistics, 37  
STD(), 22, 23, 39  
STDDEV(), 22, 23, 39  
STDDEV\_POP(), 22, 23, 39

STDDEV\_SAMP(), 22, 23, 39  
String column encoding, 38  
String functions, 21, 23, 31, 39  
SUBTIME(), 26  
SUM(), 23  
SYSDATE(), 26

## T

TABLESAMPLE, 20  
Temporal data types, 23, 31  
Temporal functions, 20, 26, 31, 34, 35, 36, 39  
Temporal type comparison, 38  
TEXT, 39  
TIME, 23, 25  
TIMEDIFF(), 26  
Timeseries forecasting, 29  
TIMESTAMP, 23, 25  
Timezone, 25  
Timezone functions, 20, 39  
TIME\_TO\_SEC(), 39  
TO\_DAYS(), 39  
TO\_SECONDS(), 39  
TRUNCATE TABLE, 29

## U

UNCOMPRESS(), 28  
UNCOMPRESSED\_LENGTH(), 28  
UNION, 39  
UNION ALL, 39  
UNIX\_TIMESTAMP(), 39  
Unload tables, 20

## V

Variable-length encoding, 38  
VARIANCE(), 22, 23, 39  
VARLEN encoding, 39  
VAR\_POP(), 22, 23, 39  
VAR\_SAMP(), 22, 23, 39  
Views, 34

## W

WEEK(), 35  
WEIGHT\_STRING, 25  
Window functions, 22, 35  
WITH ROLLUP, 35

## Y

YEAR, 23  
YEARWEEK(), 35

## Z

Zone maps, 14, 29

