

ORACLE

# Building AI Applications on your MySQL Database

---

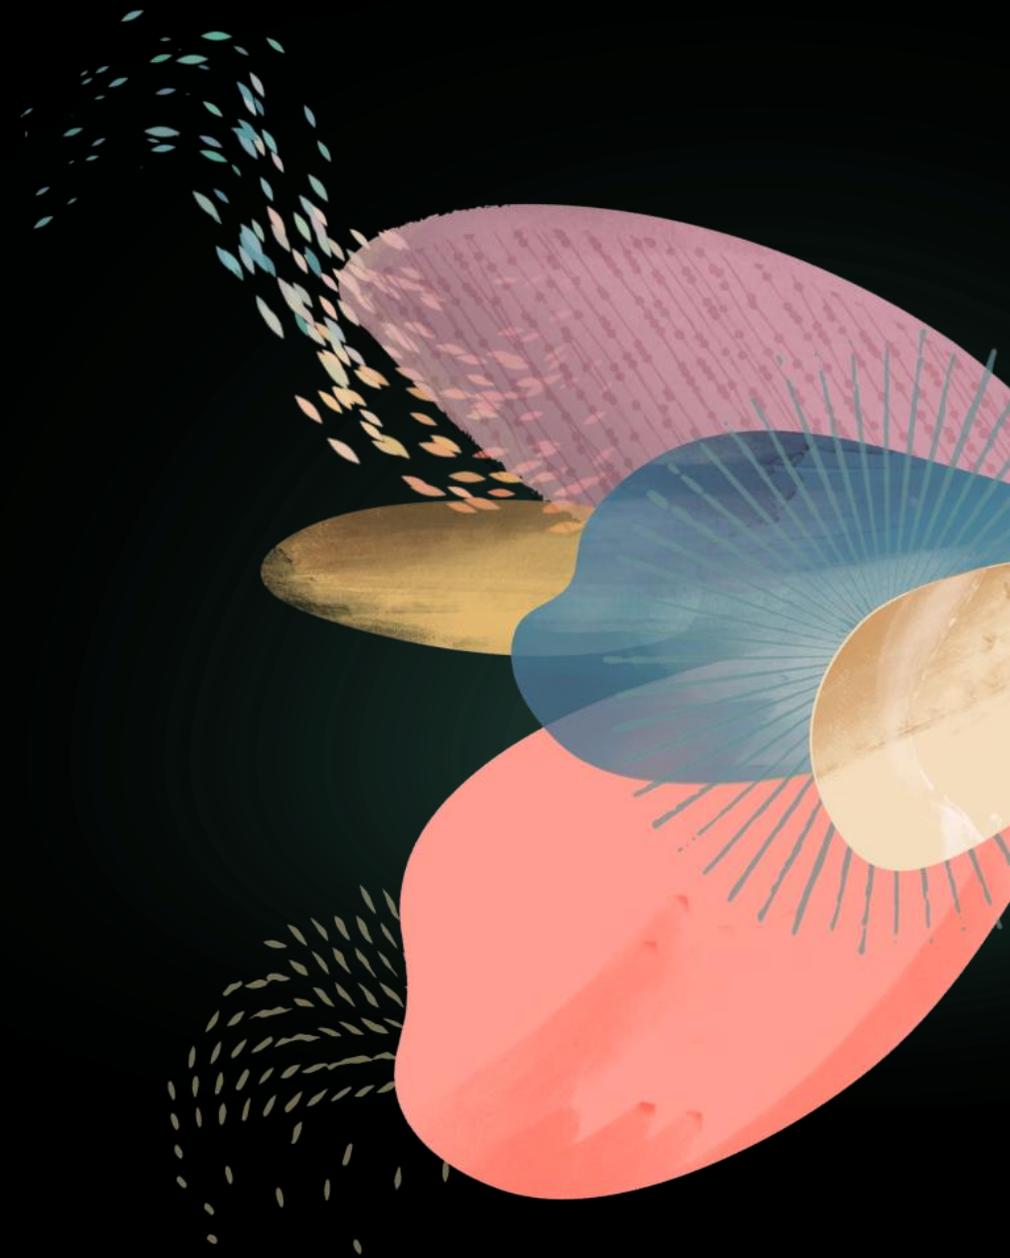
Ridha Chahed

preFOSDEM MySQL Belgian Days 2026



ORACLE

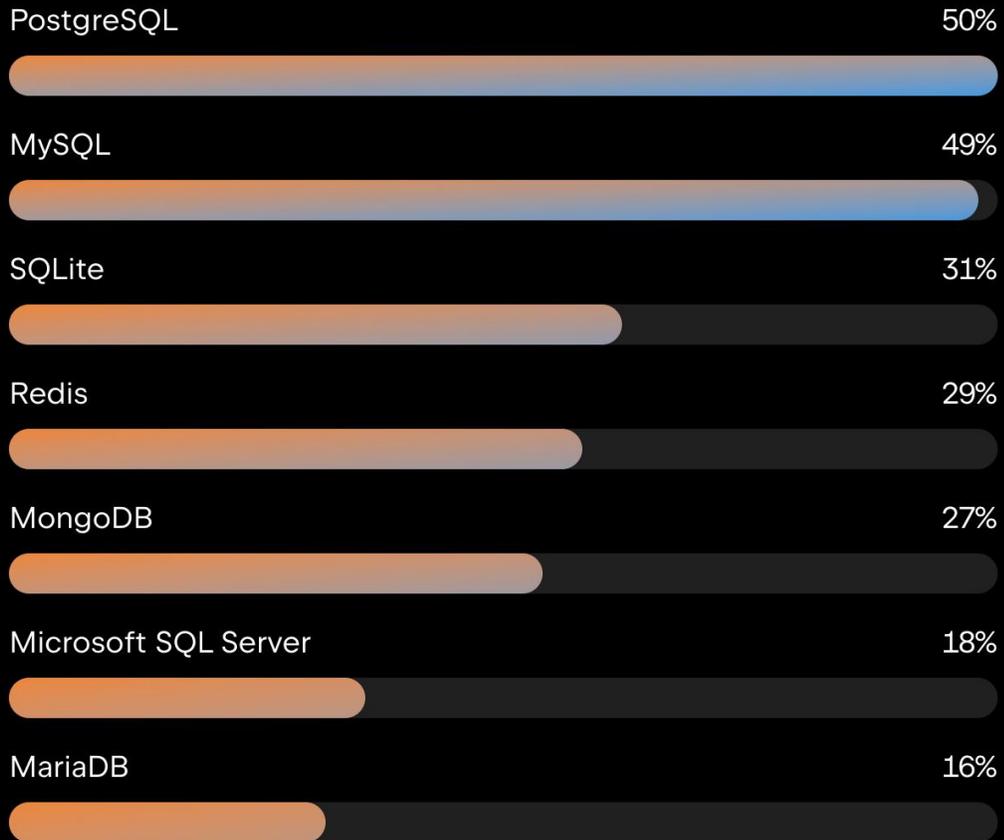
# Developer Ecosystem



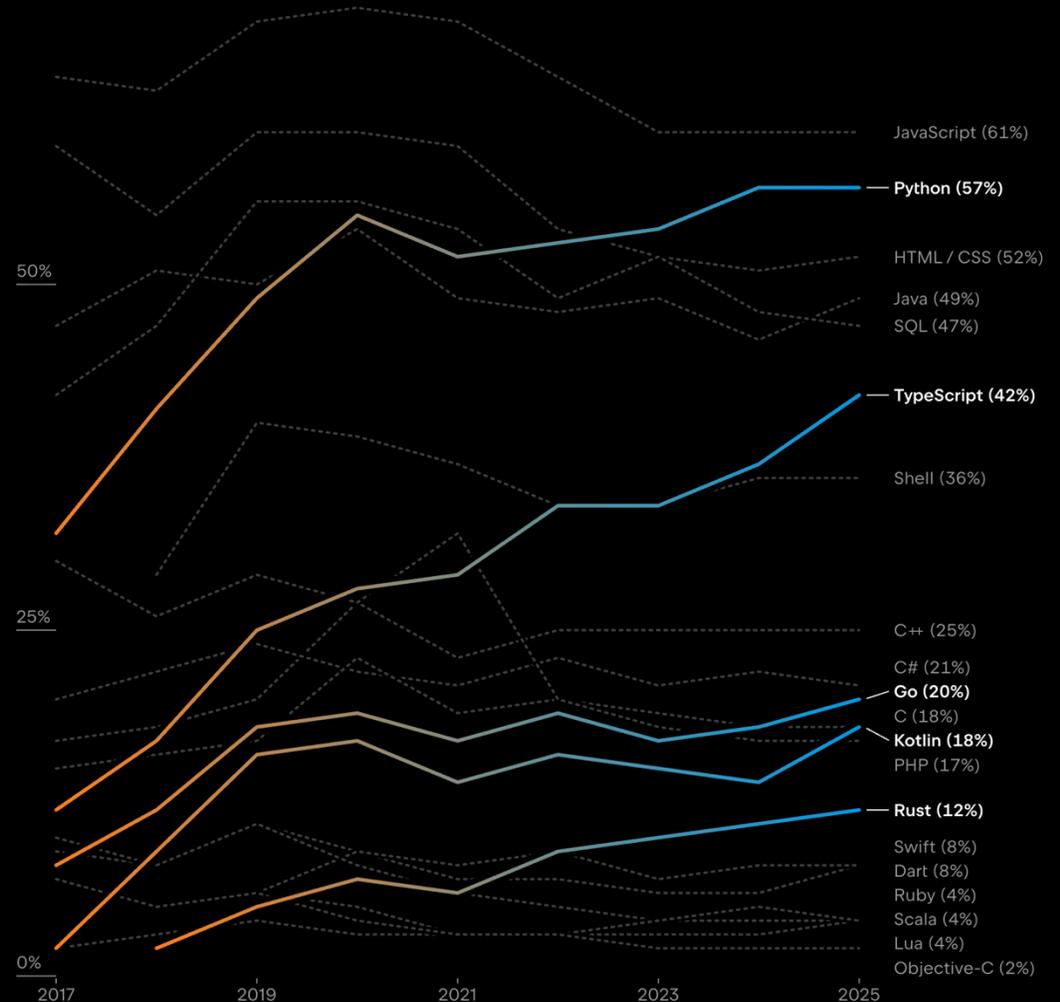
# State of the developer ecosystem 2025



## DATABASE USAGE IN 2025



## PROGRAMMING LANGUAGES



# State of the AI ecosystem 2025



Developers regularly use...

...AI tools for coding and development

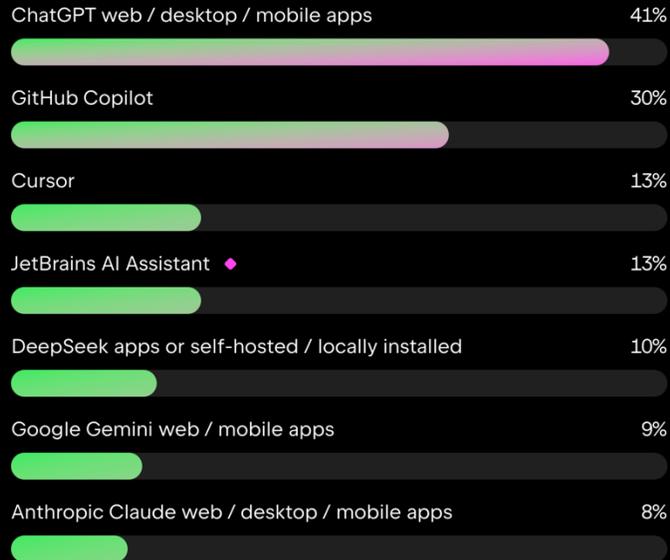
85%

...at least one AI coding assistant, agent, or code editor

62%

## AI TOOLS

regularly used for coding and development by developers



MOST LIKELY DELEGATED TO AI

## Top 5

LEAST LIKELY DELEGATED TO AI

Writing boilerplate, repetitive code

1

Communicating through email and messaging

Searching for development-related information on the internet

2

Writing application logic code

Converting code to other languages

3

Performing actions in the terminal / CLI

Writing code comments or code documentation

4

Debugging

Summarizing recent code changes

5

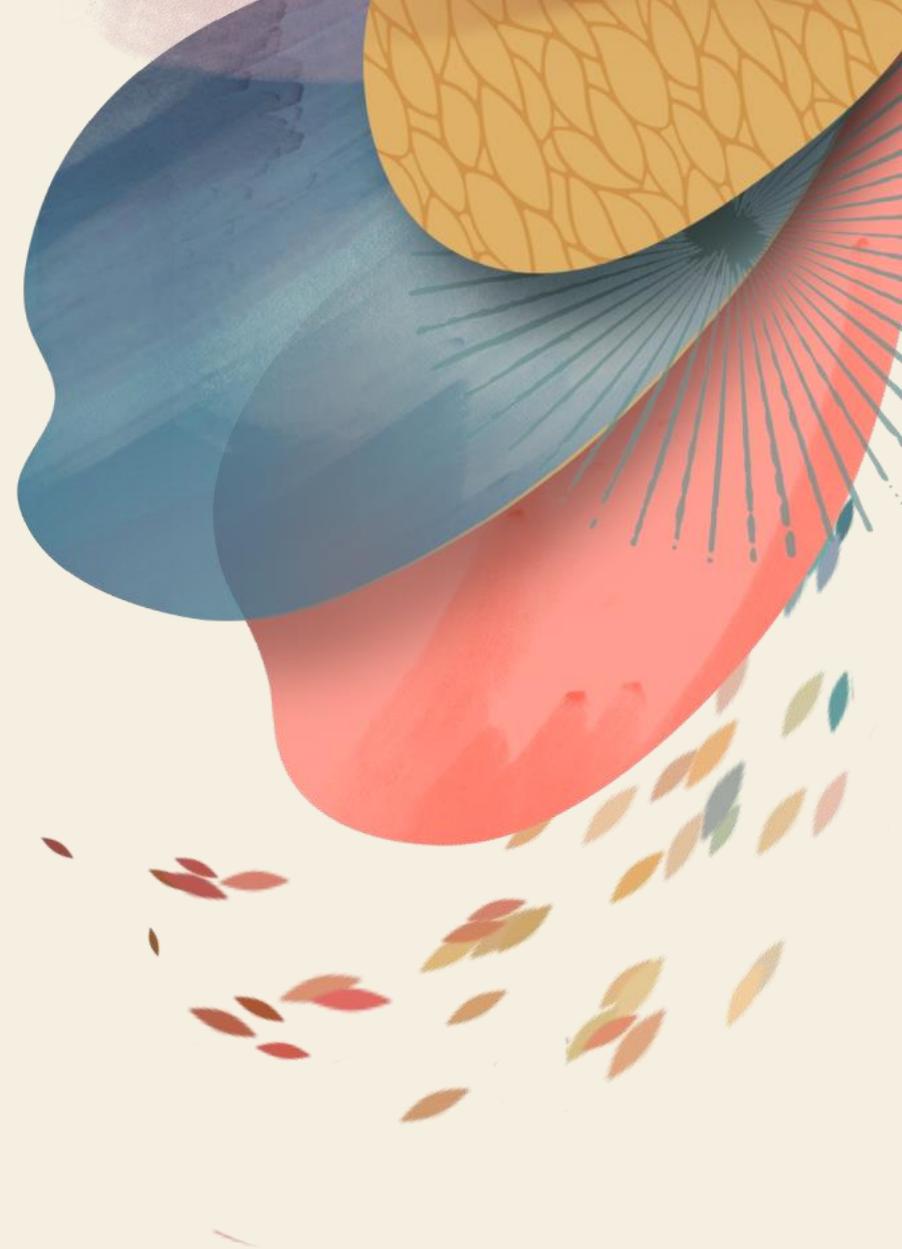
Understanding code



# MySQL MCP Server

---

Streamlining Data Innovation with Integrated Tools

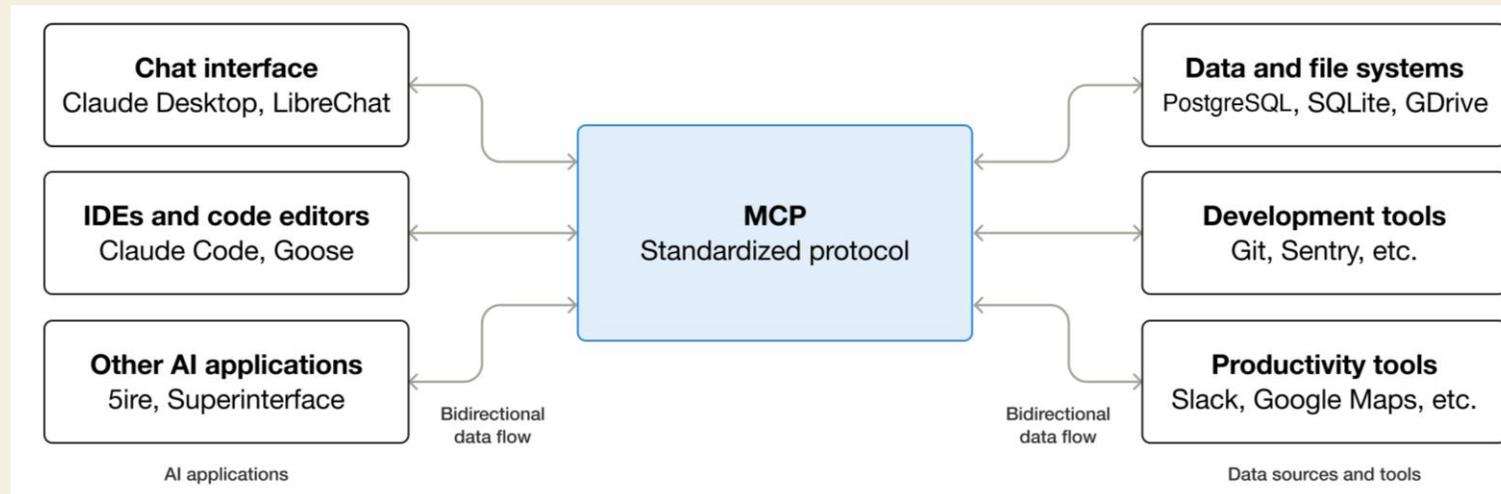


# MCP to connect your LLM to your database

MCP (Model Context Protocol) is an **open-source** standard for connecting AI applications to external systems.

AI applications like Claude or ChatGPT can connect to **data sources** like a database enabling them to access key information and perform tasks.

As an example, enterprise chatbots can connect to multiple **databases** across an organization empowering users to analyze data using chat.

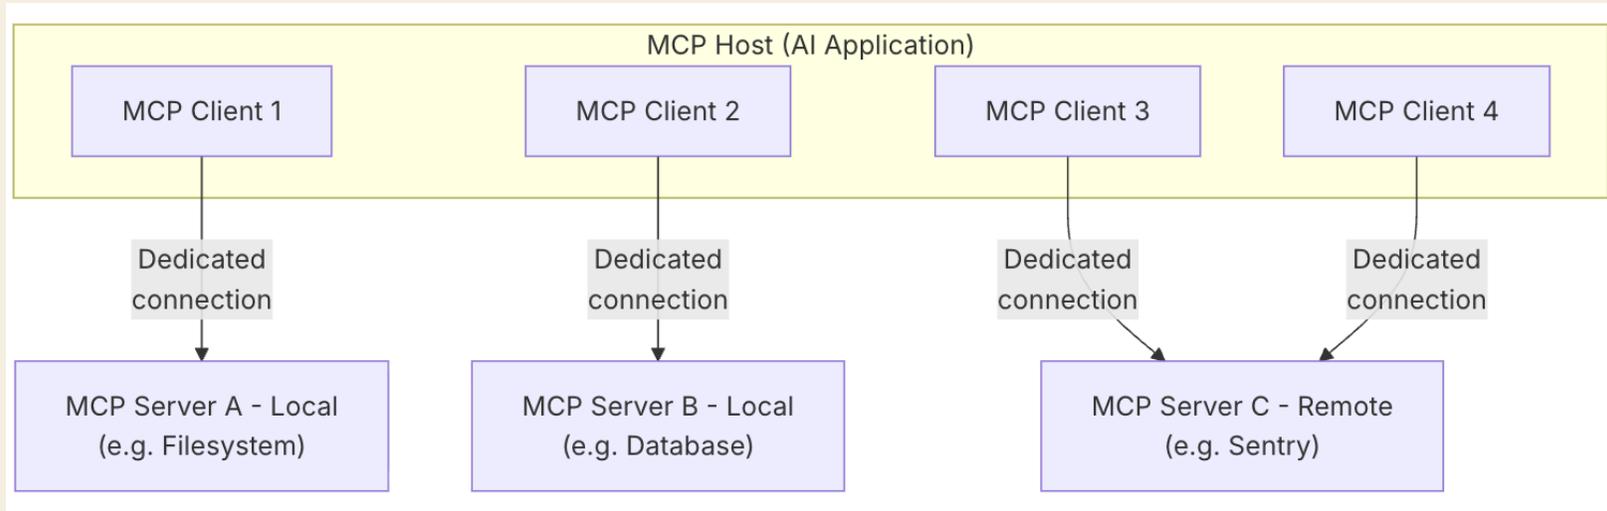


# MCP Architecture

MCP follows a **client-server** architecture where an MCP host — an AI application like Claude code or ChatGPT — establishes connections to one or more MCP servers.

The key participants in the MCP architecture are:

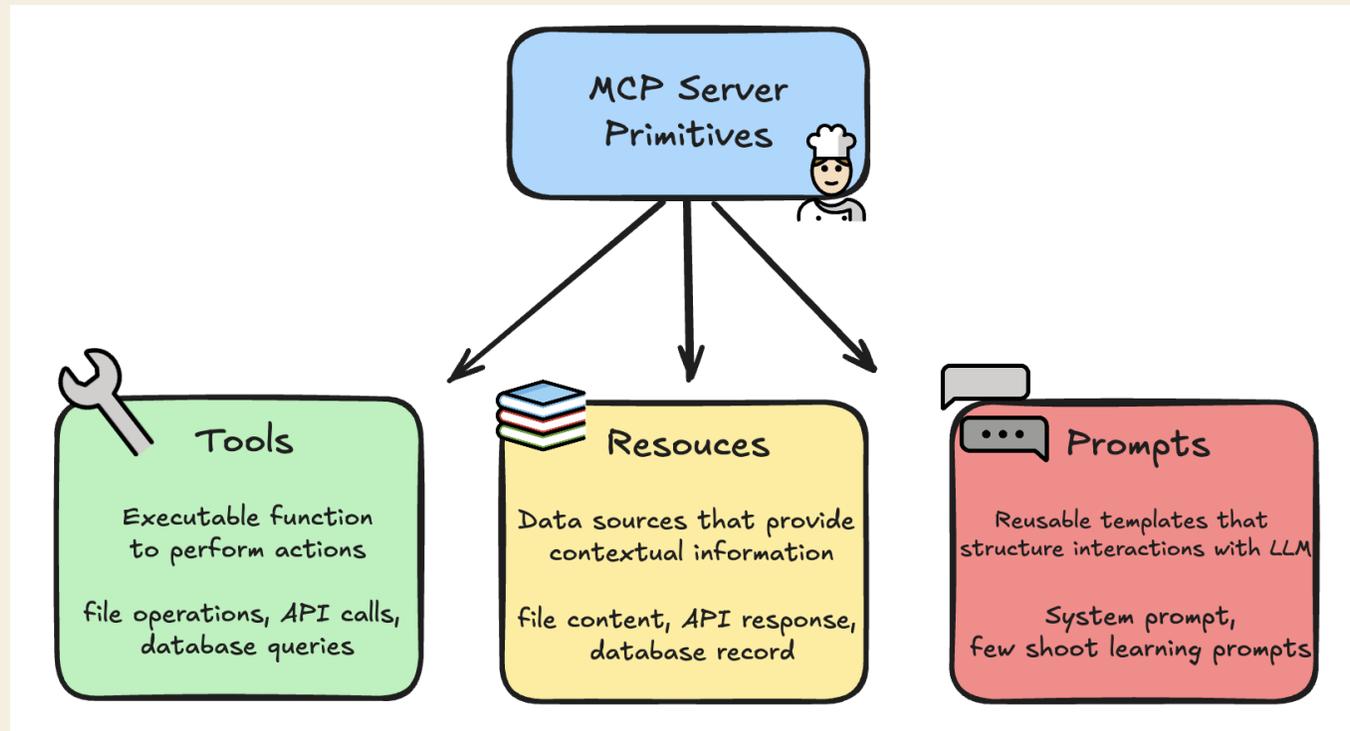
- **MCP Host:** The AI application that coordinates and manages one or multiple MCP clients
- **MCP Client:** A component that maintains a connection to an MCP server and obtains context from an MCP server for the MCP host to use
- **MCP Server:** A program that provides context to MCP clients



# MCP Primitives

MCP primitives are the most important concept within MCP. They define what clients and servers can **offer each other**.

These primitives specify the types of contextual information that can be shared with AI applications and the range of actions that can be performed.



# MySQL Open Source MCP Server



```
def _execute_sql_tool(
    connection: Union[str, MySQLConnectionAbstract],
    sql_script: str,
    params: list = None,
) -> str:
    """
    Execute a SQL script on the specified database connection.

    Args:
        connection: Union[str, MySQLConnectionAbstract]: Information defining the database connection
        to use. Allows for reusing a db connection.
        sql_script (str): The SQL statement to execute. Can be a query or a DML/DDD statement.
        params (list, optional): List of parameters to use for parameterized SQL scripts. If None,
        executes with no bind variables.

    Returns:
        str: JSON-encoded result of the query—if rows are returned, their content as a list; otherwise,
        null.
        In case of error, returns a JSON object with fields: "error", "sql_script", and "params".

    Example:
        result = _execute_sql_tool("my_conn", "SELECT * FROM users WHERE id = %s", [42])
    """
    ...

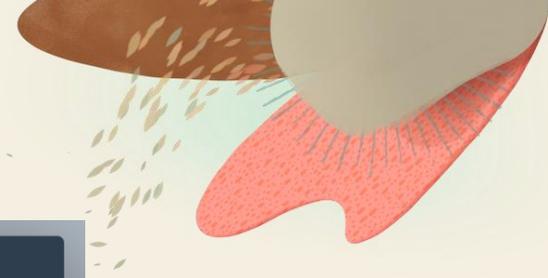
    with db_connection.cursor() as cursor:
        results = []
        cursor.execute(sql_script, params or [])

        # Read results from possibly multiple statements
        while True:
            if cursor.with_rows:
                results.extend(cursor.fetchall())

            # Move to the next result set
            if not cursor.nextset():
                break

        db_connection.commit()

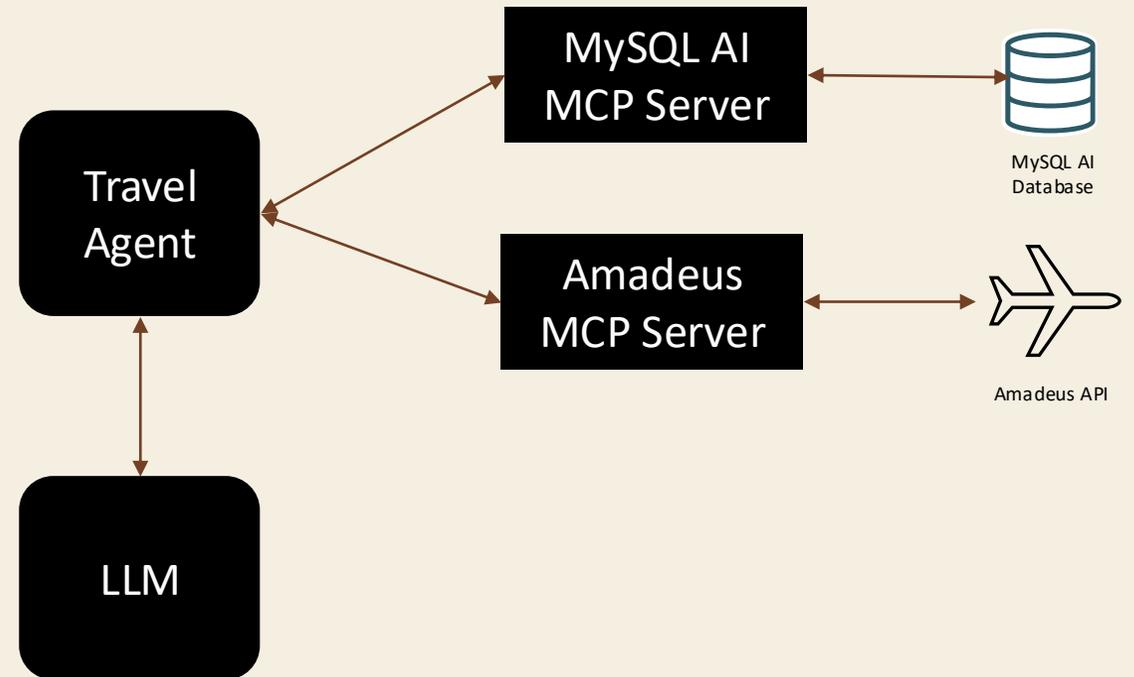
    return json.dumps(results, cls=CustomJSONEncoder)
    ...
```



# MCP Server for Travel Agent

Connects to two MCP servers

- **MySQL AI MCP Server** for booking and managing hotel reservations in a travel agency
  - `execute_sql_tool`
  - `list_tables`
  - `get_table_info`
  - ...
- **Amadeus MCP Server** for searching flights



# Travel Agency Backend

## Tables in the BOOKING database



### HOTELS

Master directory containing 10 hotels across major US cities with ratings and amenities



### ROOM\_TYPES

Inventory catalog defining 23 different room configurations with capacity and availability counts for each hotel



### ROOM\_AVAILABILITY

Real-time pricing and inventory tracking available rooms based on demand.



### BOOKINGS

Customer reservations with guest details, dates, pricing and status tracking.

## Utility procedures



### SearchHotels

Hotel search with filters for location, dates, capacity and pricing



### MakeReservation

Complete booking workflow with inventory management and pricing calculation



### CancelReservation

Cancellation processing with automatic inventory restoration

# Demo time

M List all hotels in our booking tool

D Execute sql tool

## 🏠 Complete Hotel Directory

Here are all the hotels available in our booking system:

### Boston 🏛️

**Historic Inn** (4.6 ⭐) - ID: 7

- **Amenities:** WiFi, Restaurant, Parking
- **Perfect for:** Historical exploration, business travel

### Chicago 🏙️

**Business Center Hotel** (4.1 ⭐) - ID: 4

- **Amenities:** WiFi, Gym, Business Center, Parking
- **Perfect for:** Business travelers, conferences

### Denver 🏔️

**Airport Hotel** (3.8 ⭐) - ID: 6

- **Amenities:** WiFi, Airport Shuttle, Parking, Restaurant
- **Perfect for:** Transit stays, mountain gateway

### Las Vegas 🎰

**City Center Hotel** (4.0 ⭐) - ID: 10

- **Amenities:** WiFi, Casino, Pool, Parking
- **Perfect for:** Entertainment, gaming, shows

M I have a new customer:

John Doe  
john.doe@oracle.com  
+1 123 1234567

Please book a room for Mr. Doe and his spouse in Chicago. They will attend a conference from November 3 to November 7, 2025.

Perfect! I'll book a room for John Doe and his spouse at the Business Center Hotel in Chicago - it's ideal for conference attendees with its business amenities. Let me make the reservation for November 3-7, 2025.

D Execute sql tool

Great! I can see there are two room types available at the Business Center Hotel in Chicago. Let me book the Executive room (room\_type\_id: 10) for John Doe and his spouse, as it's likely to be more suitable for conference attendees:

D Execute sql tool

🎉 **BOOKING CONFIRMED FOR JOHN DOE!**

### Reservation Details for John Doe

**Booking ID:** 520

**Status:** SUCCESS ✅

📍 **Destination:** Chicago, IL

**Hotel:** Business Center Hotel (4.1 ⭐)

**Room Type:** Executive Room (capacity: 2 people)

**Check-in:** November 3, 2025

**Check-out:** November 7, 2025

**Duration:** 4 nights

# Demo time

M Oh, they would like to stay a day longer to do some sightseeing. Can you extend their reservation?

Of course! I'll extend John Doe's Chicago reservation by one day so they can enjoy some sightseeing. Let me cancel their current booking and create a new one for November 3-8, 2025 (5 nights instead of 4).

D Execute sql tool

Perfect! Now let me book the extended stay for November 3-8, 2025 (5 nights) so they can enjoy an extra day of sightseeing:

D Execute sql tool

**🎉 RESERVATION EXTENDED SUCCESSFULLY!**

## Updated Reservation Details for John Doe

**New Booking ID:** 521

**Status:** SUCCESS 

**📍 Destination:** Chicago, IL

**Hotel:** Business Center Hotel (4.1 )

**Room Type:** Executive Room (capacity: 2 people)

**Check-in:** November 3, 2025

**Check-out:** November 8, 2025

**Duration:** 5 nights (Extended from 4 nights!)

M List a few flights from New York to Chicago that could work for the Doe family.

I'll search for flights from New York to Chicago for John Doe and his spouse. Since they need to arrive for their conference on November 3, 2025, let me look for flights departing November 2 or early November 3, and returning November 8, 2025.

A Search flight offers

Excellent! I found several great flight options for John Doe and his spouse traveling from New York (LaGuardia) to Chicago for their conference and extended sightseeing trip. Here are the best options:

### **Flight Options: New York (LGA) Chicago (ORD)**

**For John Doe & Spouse (2 passengers) | November 3-8, 2025**

#### **BEST VALUE - EARLY ARRIVAL**

**Option 1: United Airlines - \$579.94 total  RECOMMENDED**

**Outbound:** November 3, 2025

- **LGA → ORD** (Non-stop)
- Depart: **6:00 AM** → Arrive: **7:25 AM** (local time)
- Flight Duration: 2h 25m
- **Perfect timing:** Arrive early for conference setup

**Return:** November 8, 2025

- **ORD → LGA** (Non-stop)
- Depart: **1:00 PM** → Arrive: **4:05 PM**
- Flight Duration: 2h 5m
- **Perfect timing:** Full morning for sightseeing

#### **CONVENIENT TIMING OPTIONS**

**Option 2: United Airlines - \$643.86 total**

Reply to Claude...

+   Research

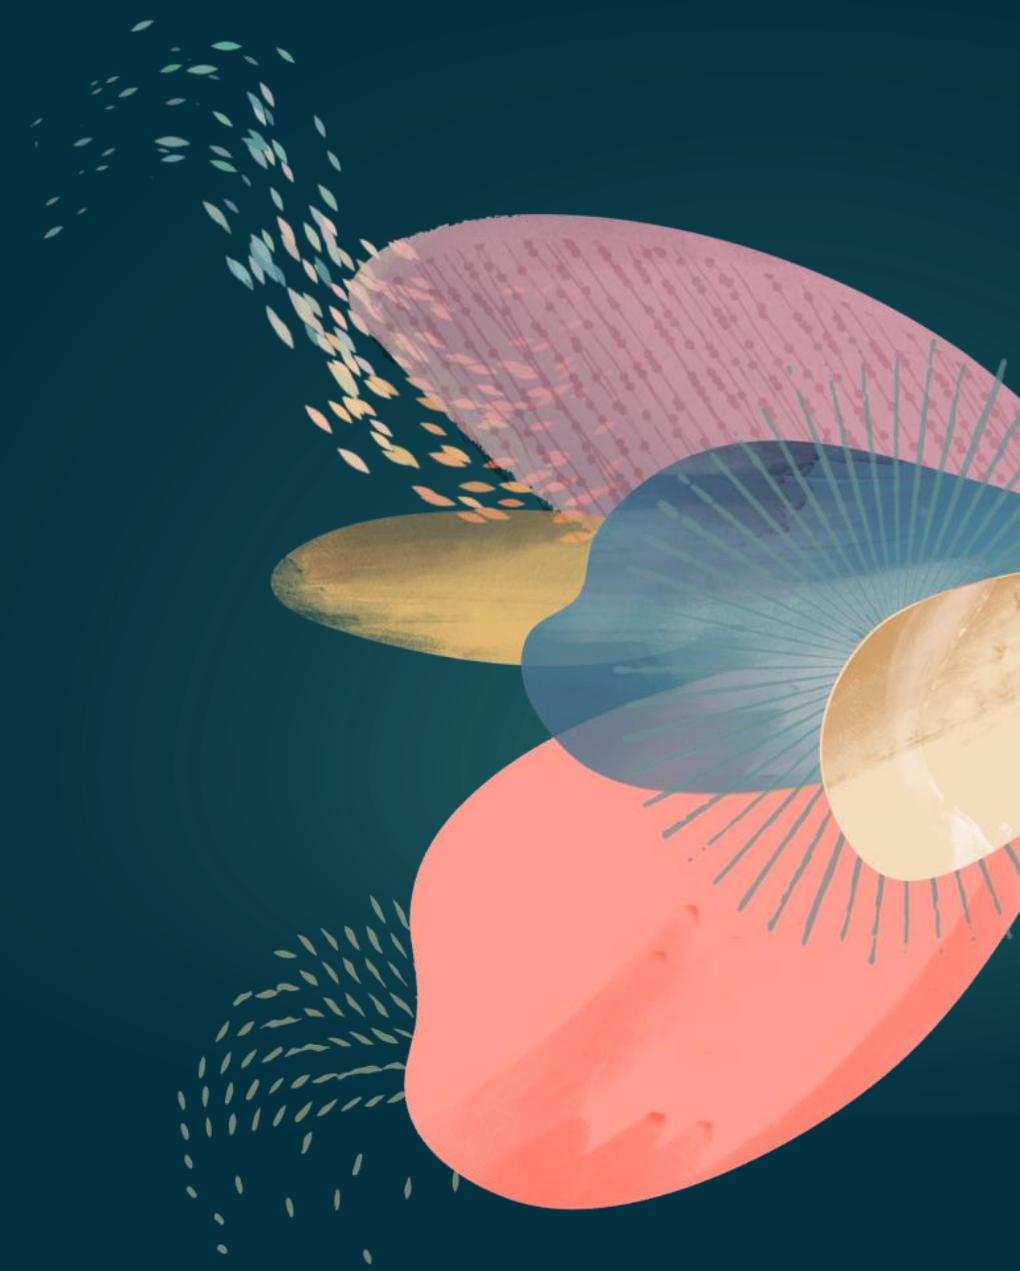
Claude Sonnet 4 



ORACLE

# MySQL Studio

Streamlining Data Innovation with Integrated Tools



# MySQL Studio

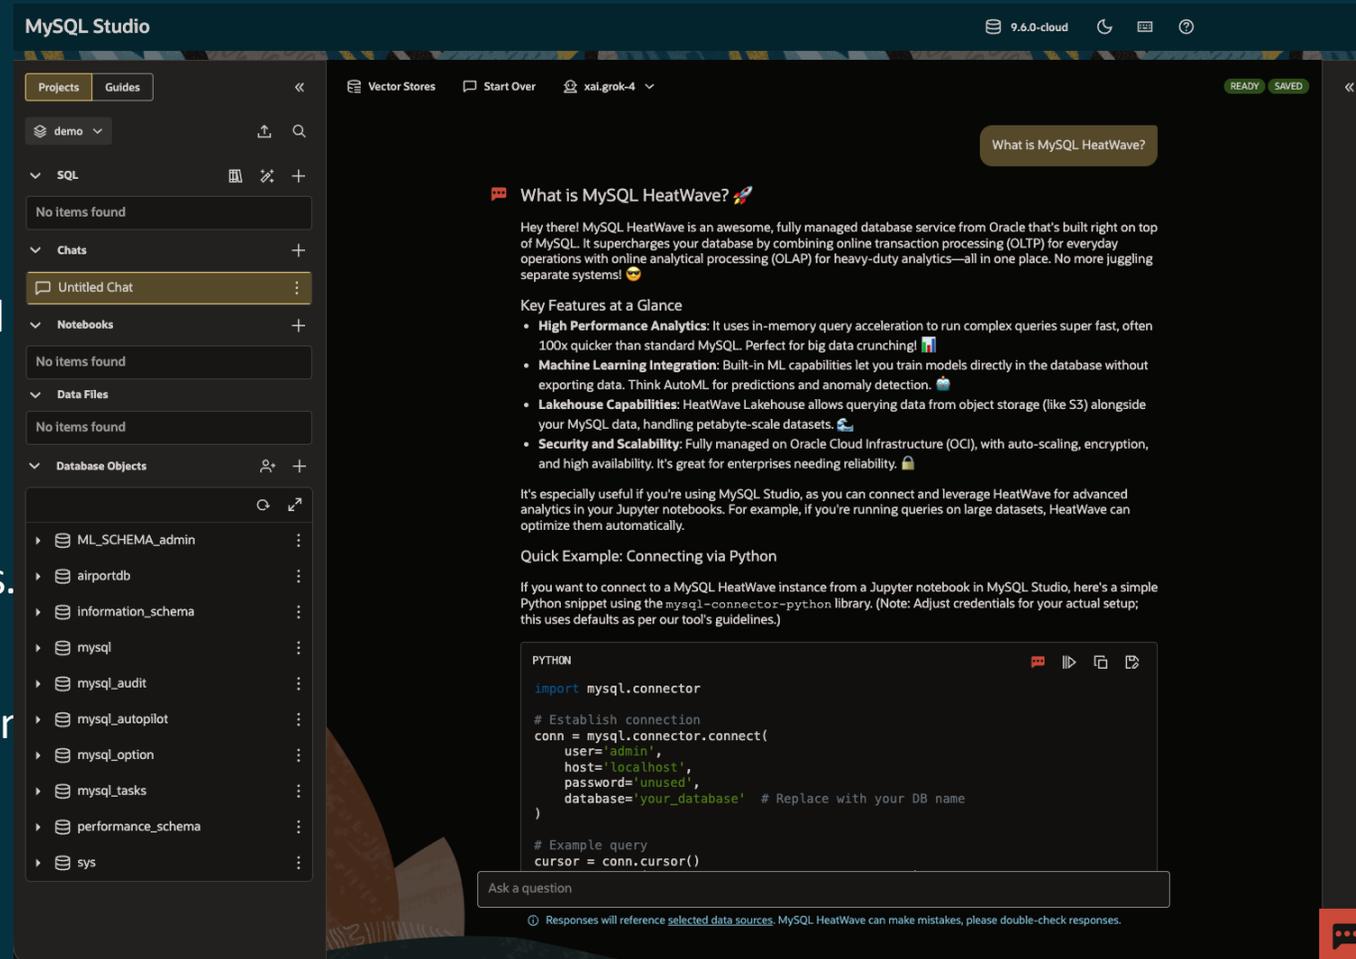
## Streamlining Data Innovation with Integrated Tools

**Integrated Environment** – Everything (SQL, AI Chats and notebooks) runs in OCI; no installs, plugins, or setups.

**Jupyter Compatible** – Import existing notebooks and run Python and SQL natively.

**Project Organization** – Manages SQL scripts, AI interactions, and data files cohesively within projects.

**AI Assistance** - Natural language queries, explanation which speeds up everyday work.



# MySQL Studio

## Lowering the Barriers

- Simplifies migration from on-prem to cloud (same console as MySQL AI)
- No need to setup and bring other tools to use the database
- Curating collection of built-in *notebook guides* to get started quickly with advance data-plane features
- Integrated SQL, AI assistance, and visualization reduce tool sprawl and context switching
- Enables organizations to quickly unlock analytics and performance benefits

## Supported Databases

- MySQL 8.0.41 -> 9.6.0 (and beyond)

## Supported deployments

- > 9.4.1 needed for Generative AI
- HeatWave cluster needed for:
  - Generative AI
  - AutoML
  - NL2SQL
  - AI Assistance



# SQL

- Database Objects Explorer
  - Quickly view schemas, view, procedures...
- Built in snippets to accomplish tasks and learn about
  - Create database, users and tables
  - View sample from a table
- NL2SQL to explore data quickly
- Visualize the results
- Query history
- Async task viewer / explorer
- Explain queries

The screenshot shows the MySQL Studio interface. On the left is the Database Objects Explorer showing a tree view of the 'airportdb' database with tables like 'airline', 'airplane', 'airplane\_type', 'airport', 'airport\_geo', and 'booking'. The main window displays a SQL query that counts flights by aircraft type. The query results are shown as a pie chart with 13 segments, each representing an aircraft type and its corresponding number of flights.

```
1 SELECT
2   at.identifier AS `Aircraft Name`,
3   COUNT(f.flight_id) AS `Number of Flights`
4 FROM
5   airportdb.airplane_type at
6 JOIN
7   airportdb.airplane a ON at.type_id = a.type_id
8 JOIN
9   airportdb.flight f ON a.airplane_id = f.airplane_id
10 GROUP BY
11   at.type_id, at.identifier
12 ORDER BY
13   `Number of Flights` DESC;
```

Aircraft Type	Number of Flights
Boeing 737	1
Bombardier Q Series	1
Douglas DC-9	1
Airbus A380	1
Fokker 100	1
McDonnell Douglas DC-10	1
Airbus A330	1
Embraer-ERJ-145-Familie	1
Boeing 747	1
Airbus-A320-Familie	1
Fokker 70	1
Boeing 767	1
Boeing 777	1



# AI Chat

**MySQL Studio** 9.6.0-cloud

Vector Stores Start Over xai.grok-4

READY SAVED

What is MySQL HeatWave?

**What is MySQL HeatWave?**

Hey there! MySQL HeatWave is an awesome, fully managed database service from Oracle that's built right on top of MySQL. It supercharges your database by combining online transaction processing (OLTP) for everyday operations with online analytical processing (OLAP) for heavy-duty analytics—all in one place. No more juggling separate systems! 😊

**Key Features at a Glance**

- **High Performance Analytics:** It uses in-memory query acceleration to run complex queries super fast, often 100x quicker than standard MySQL. Perfect for big data crunching! 🚀
- **Machine Learning Integration:** Built-in ML capabilities let you train models directly in the database without exporting data. Think AutoML for predictions and anomaly detection. 🧠
- **Lakehouse Capabilities:** HeatWave Lakehouse allows querying data from object storage (like S3) alongside your MySQL data, handling petabyte-scale datasets. 🌊
- **Security and Scalability:** Fully managed on Oracle Cloud Infrastructure (OCI), with auto-scaling, encryption, and high availability. It's great for enterprises needing reliability. 🛡️

It's especially useful if you're using MySQL Studio, as you can connect and leverage HeatWave for advanced analytics in your Jupyter notebooks. For example, if you're running queries on large datasets, HeatWave can optimize them automatically.

**Quick Example: Connecting via Python**

If you want to connect to a MySQL HeatWave instance from a Jupyter notebook in MySQL Studio, here's a simple Python snippet using the `mysql-connector-python` library. (Note: Adjust credentials for your actual setup; this uses defaults as per our tool's guidelines.)

```
PYTHON
import mysql.connector

# Establish connection
conn = mysql.connector.connect(
    user='admin',
    host='localhost',
    password='unused',
    database='your_database' # Replace with your DB name
)

# Example query
cursor = conn.cursor()
```

Ask a question

Responses will reference [selected data sources](#). MySQL HeatWave can make mistakes, please double-check responses.

- Chat with unstructured data
  - RAG
- Filter which vector stores to query
- View citations
- Save chats for later
- Import / export chats

# Python Notebooks

- Jupyter compatible
  - Customers can bring their own notebooks
- Set of common ML/DS libraries
  - Pandas, numpy, Ibis and MySQL connector
- IPywidgets / IPeCharts
  - Build mini apps right in the console
  - Expose full Apache ECharts library
- Curated Notebook Guide
  - Comprehensive set of guides
  - Building mini apps / chatbots

### Penguin Classifier

1. Select a model you want to use:

Model: RandomForest

2. Set penguin measurements:

Beak Length (mm): 45      Beak Depth (mm): 17  
Body Mass (g): 3500      Flipper (mm): 200

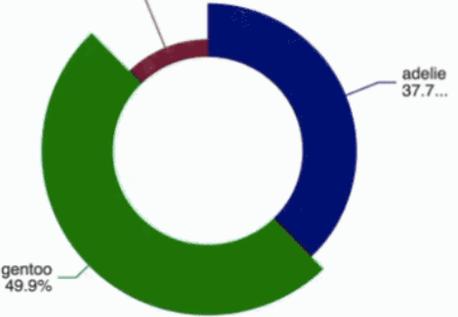
3. Predict Penguin Species

**Predicted Species: Gentoo**  
Likelihoods: Adelle: 37.8%, Gentoo: 49.9%, Chinstrap: 12.3%



- adelle
- gentoo
- chinstrap

#### Species Probability



Species	Probability
Adelle	37.7%
Gentoo	49.9%
Chinstrap	12.33%

Copyright © 2025, Oracle and/or its affiliates. All rights reserved.

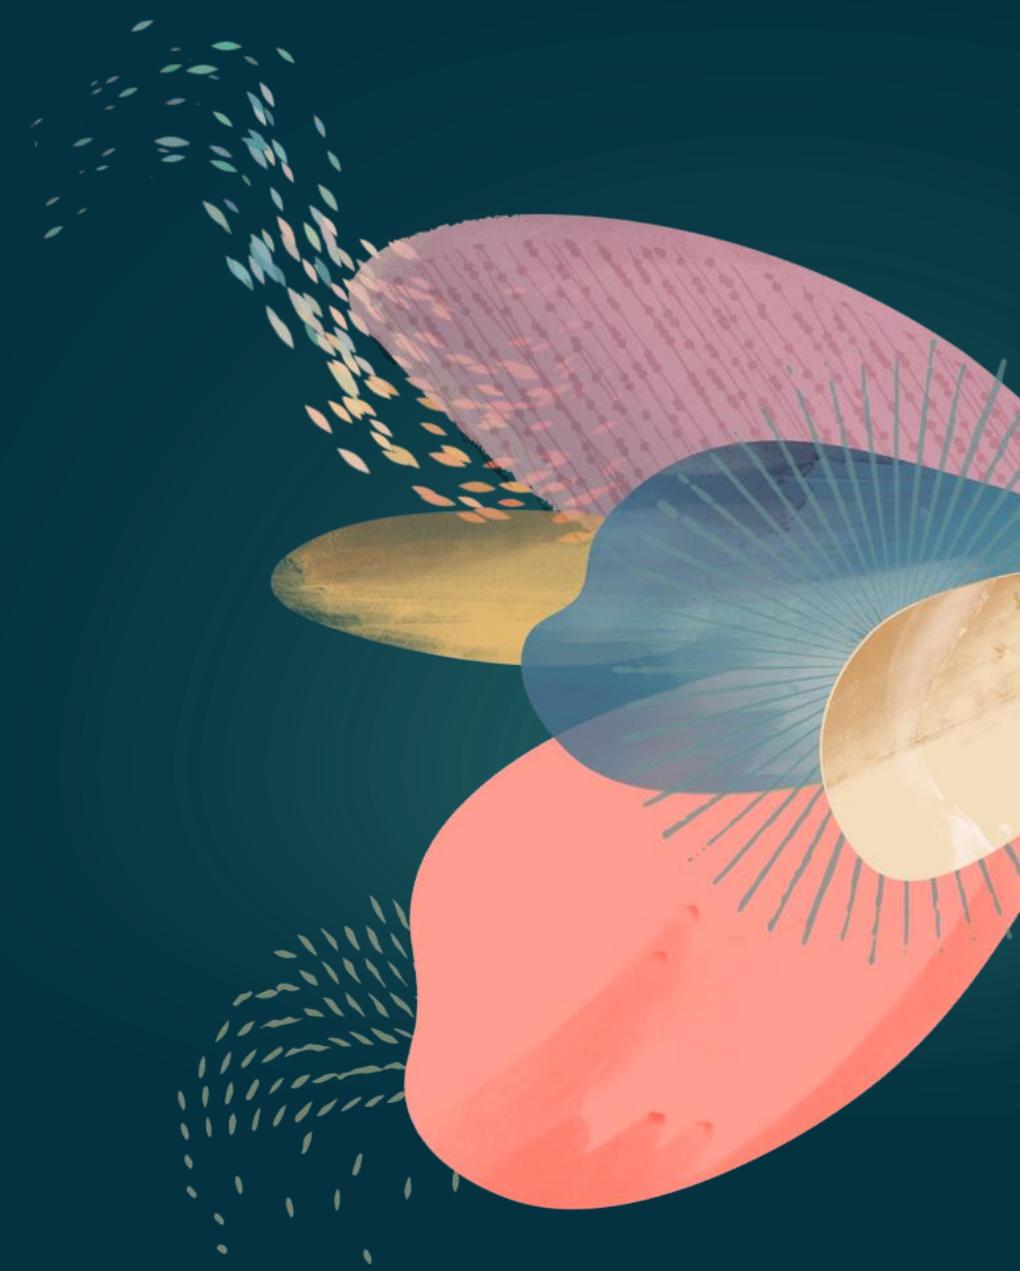


ORACLE

# MySQL AI

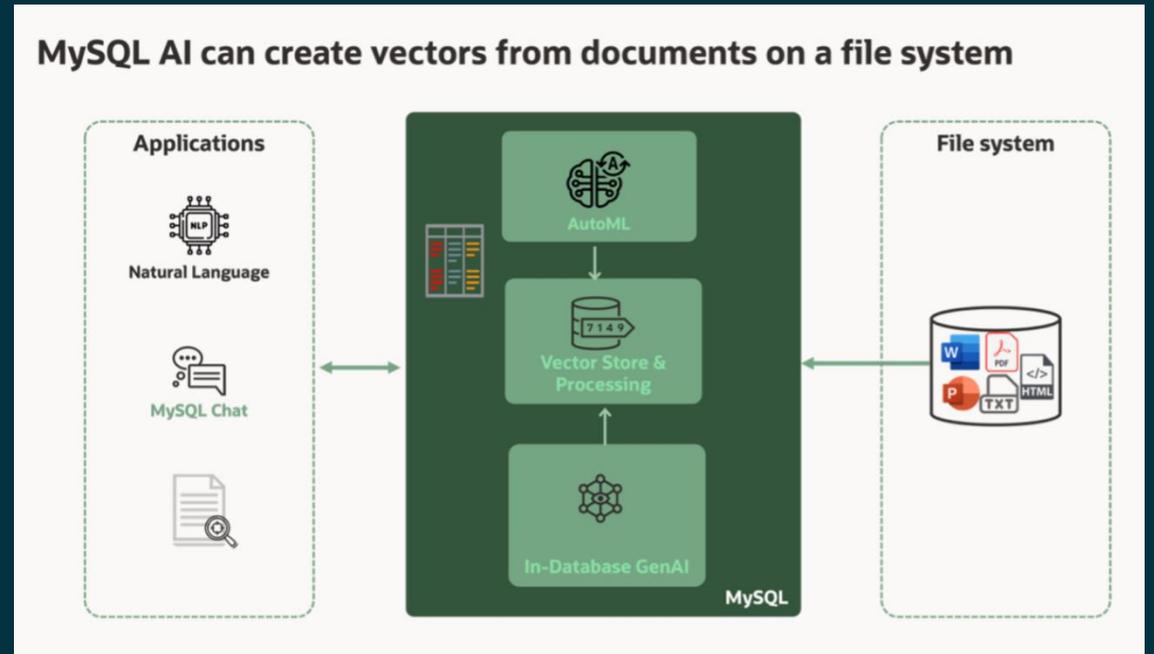
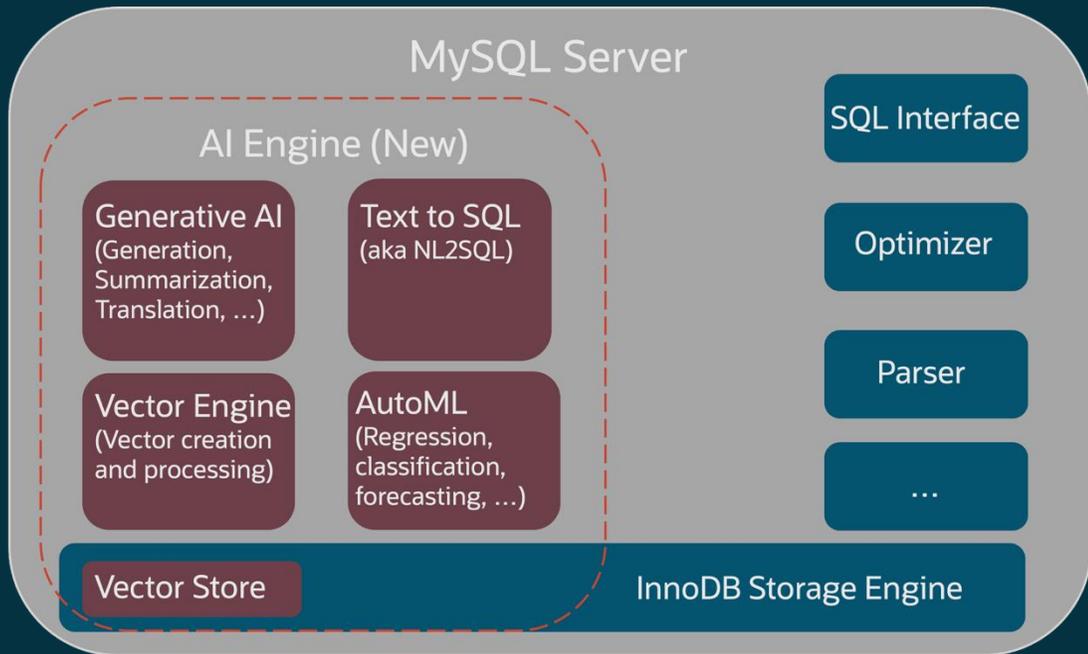
Making the MySQL database AI native

---

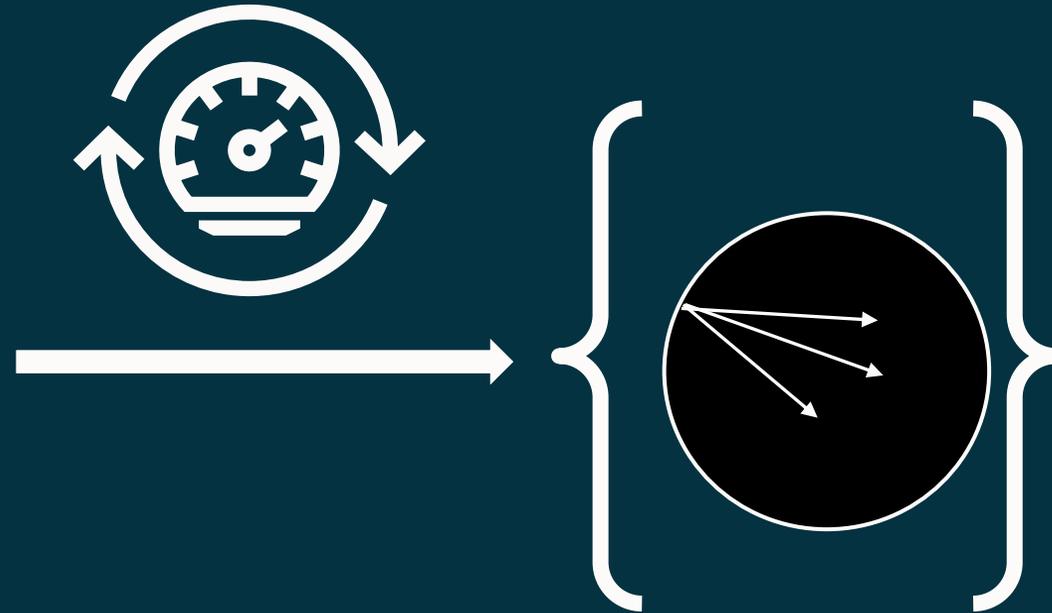


# MySQL AI

Option of **MySQL Enterprise Edition** and provides built-in predictive and generative AI capabilities, including automated machine learning, in built LLMs, and a vector store.



# Vectors or Embeddings



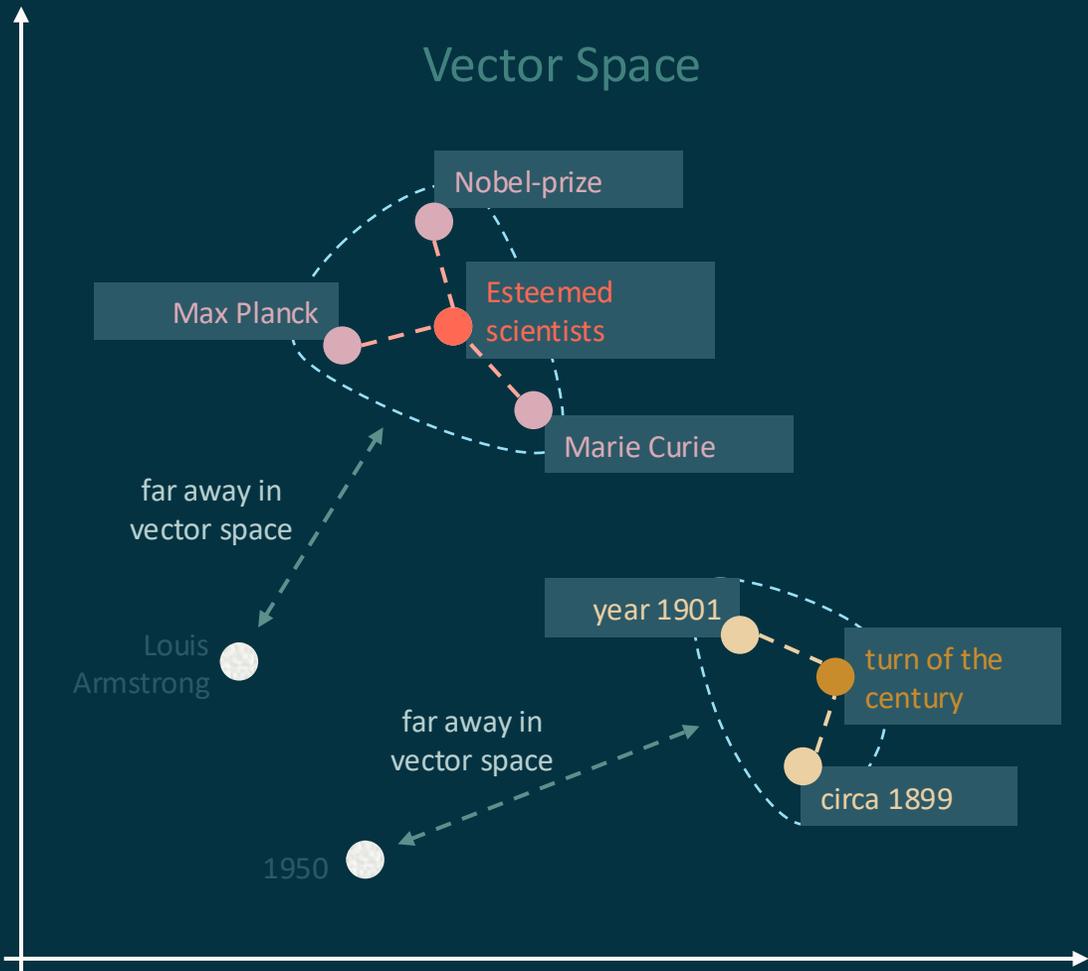
Data of any kind – images, text, documents, audio ...

- LM is used to generate VECTORS corresponding to user data
- A VECTOR an array of numbers is a coordinate in multi-dimensional space
- Finding mathematical distance between two VECTORS is similarity search





# Distance function enables semantic search



- Documents with articles about physics experiments are stored in HeatWave Vector Store (info mapped to Vector Space)
- Each point represents vector embedding of data. Data similar in semantic meaning are closer in Vector Space
- DISTANCE function measures the distance between points in the Vector Space
- Enables searching data semantically, instead of traditional SQL search which directly compares data values
- E.g. “Esteemed scientist” search key word is contextually close to Max Planck or Marie Curie



# New vector datatype in HeatWave & in MySQL

## VECTOR is a first-class citizen now



Vector as  
first-class  
data type

```
mysql> CREATE TABLE wikipedia (  
    title VARCHAR(1024),  
    page_data TEXT,  
    page_url TEXT,  
    page_embedding VECTOR(1024));
```

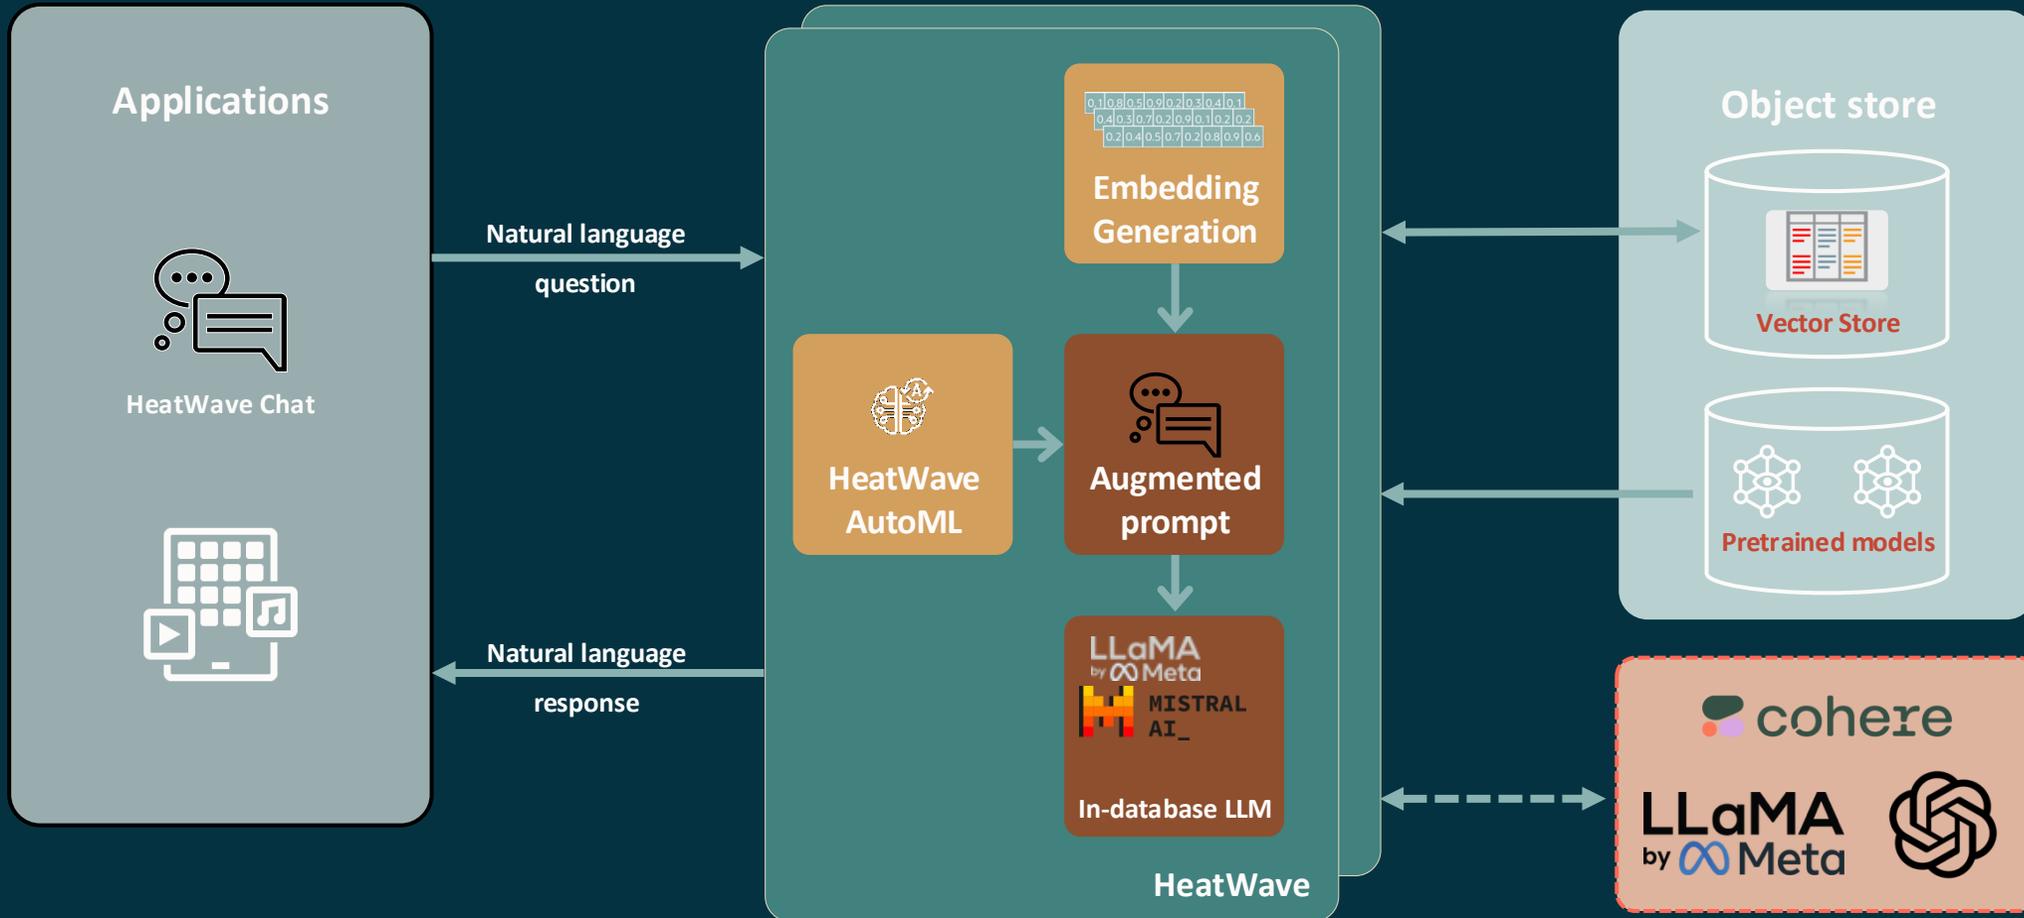
MySQL query  
syntax

```
mysql> SELECT page_url,  
    DISTANCE(page_embedding,  
             @query_embedding, "COSINE")  
    as distance  
FROM wikipedia  
ORDER by distance DESC LIMIT 10;
```

New distance function for similarity search

- L2/EUCLIDIAN
- COSINE
- DOT

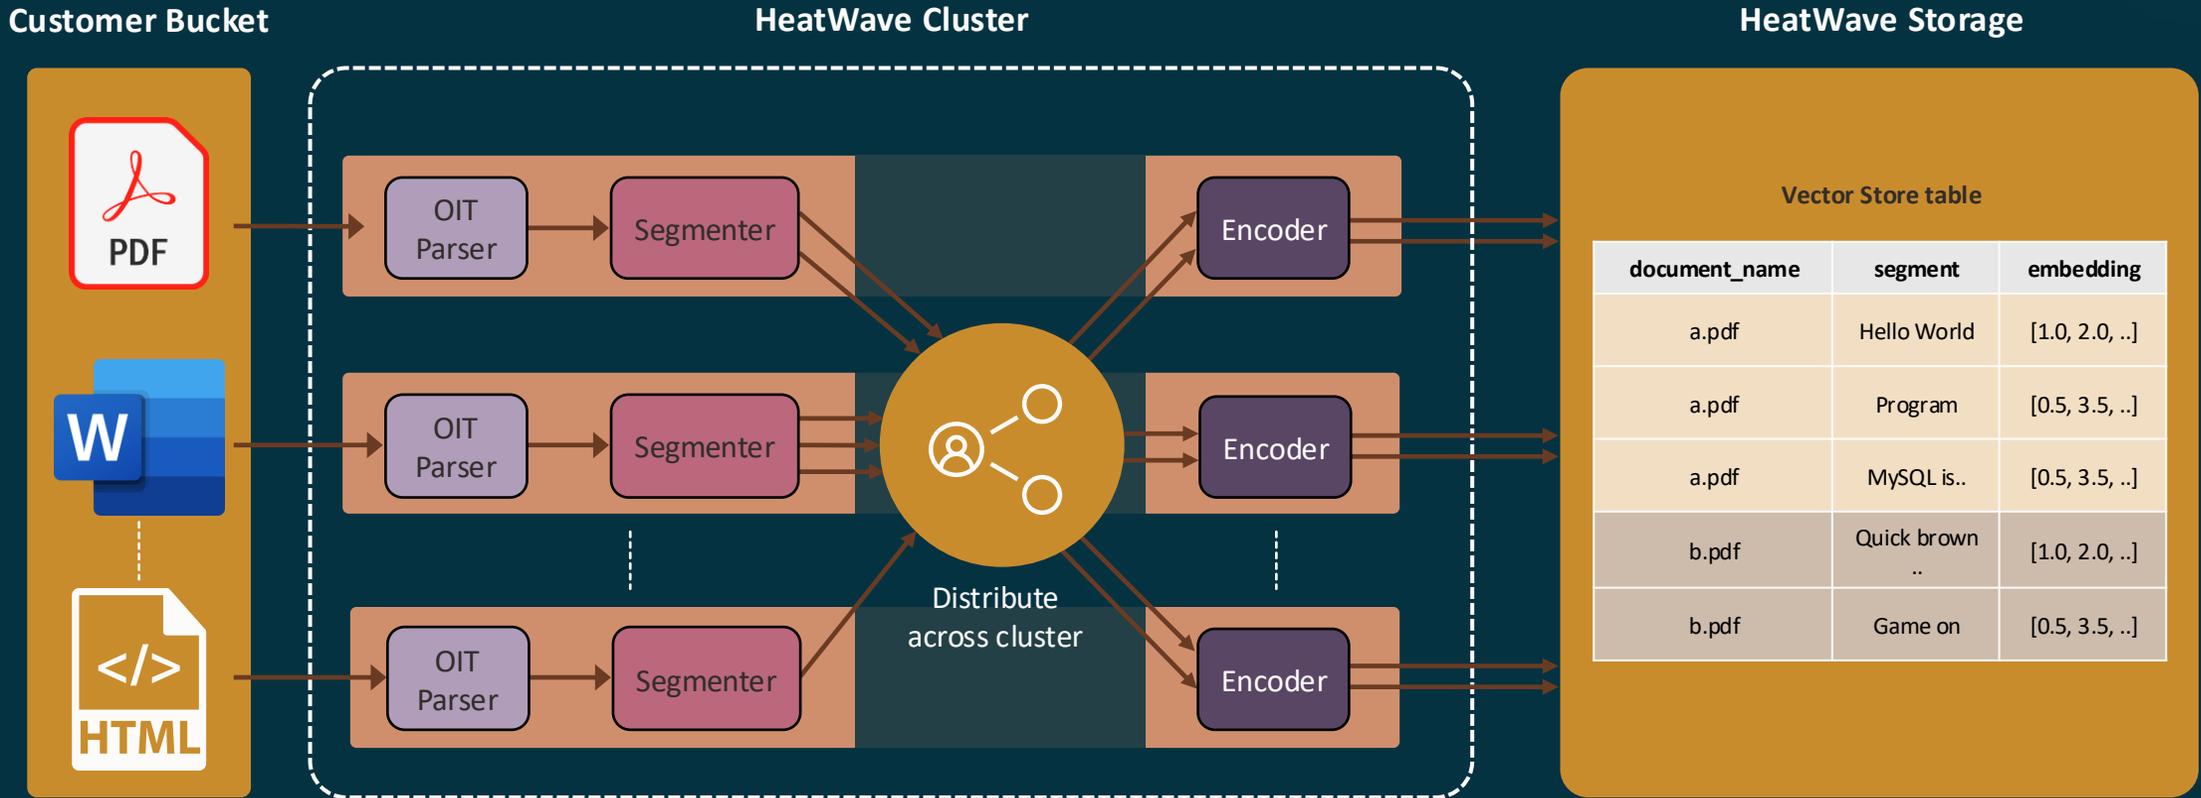
# MySQL AI on the cloud



```
> CALL sys.ML_RAG("What is the reimbursement process?", @NL_response,  
@optional_search_parameters);
```

# Create a vector store

All system resources are optimized by HeatWave



Faster than generating vector store at the application layer



# Demo

---

