

Application Development with MySQL HeatWave Using Oracle Analytics Cloud, APEX, OCI GoldenGate, and Notebooks

Salil Pradhan
Principal Product Manager
MySQL HeatWave
March 23, 2023






MySQL: An extremely popular database




MySQL is the most popular database for developers

Most popular databases



[Stackoverflow survey](#)

Which databases have you used in the last 12 months?



[Jetbrains survey](#)

Innovative organizations across many industries run MySQL

Social

facebook



LinkedIn



Pinterest

E-Commerce

Booking.com

NETFLIX

UBER



淘宝网
Taobao.com

阿里巴巴
Alibaba.com

Tech



GitHub

HubSpot

zendesk



Finance



J.P.Morgan

citi



VISA



Manufacturing

TESLA



TOYOTA



Common MySQL use cases



Social applications

Fraud detection

Authentication systems

E-Commerce

Digital payments

Content management

Retail POS systems

IoT monitoring systems

Digital marketing

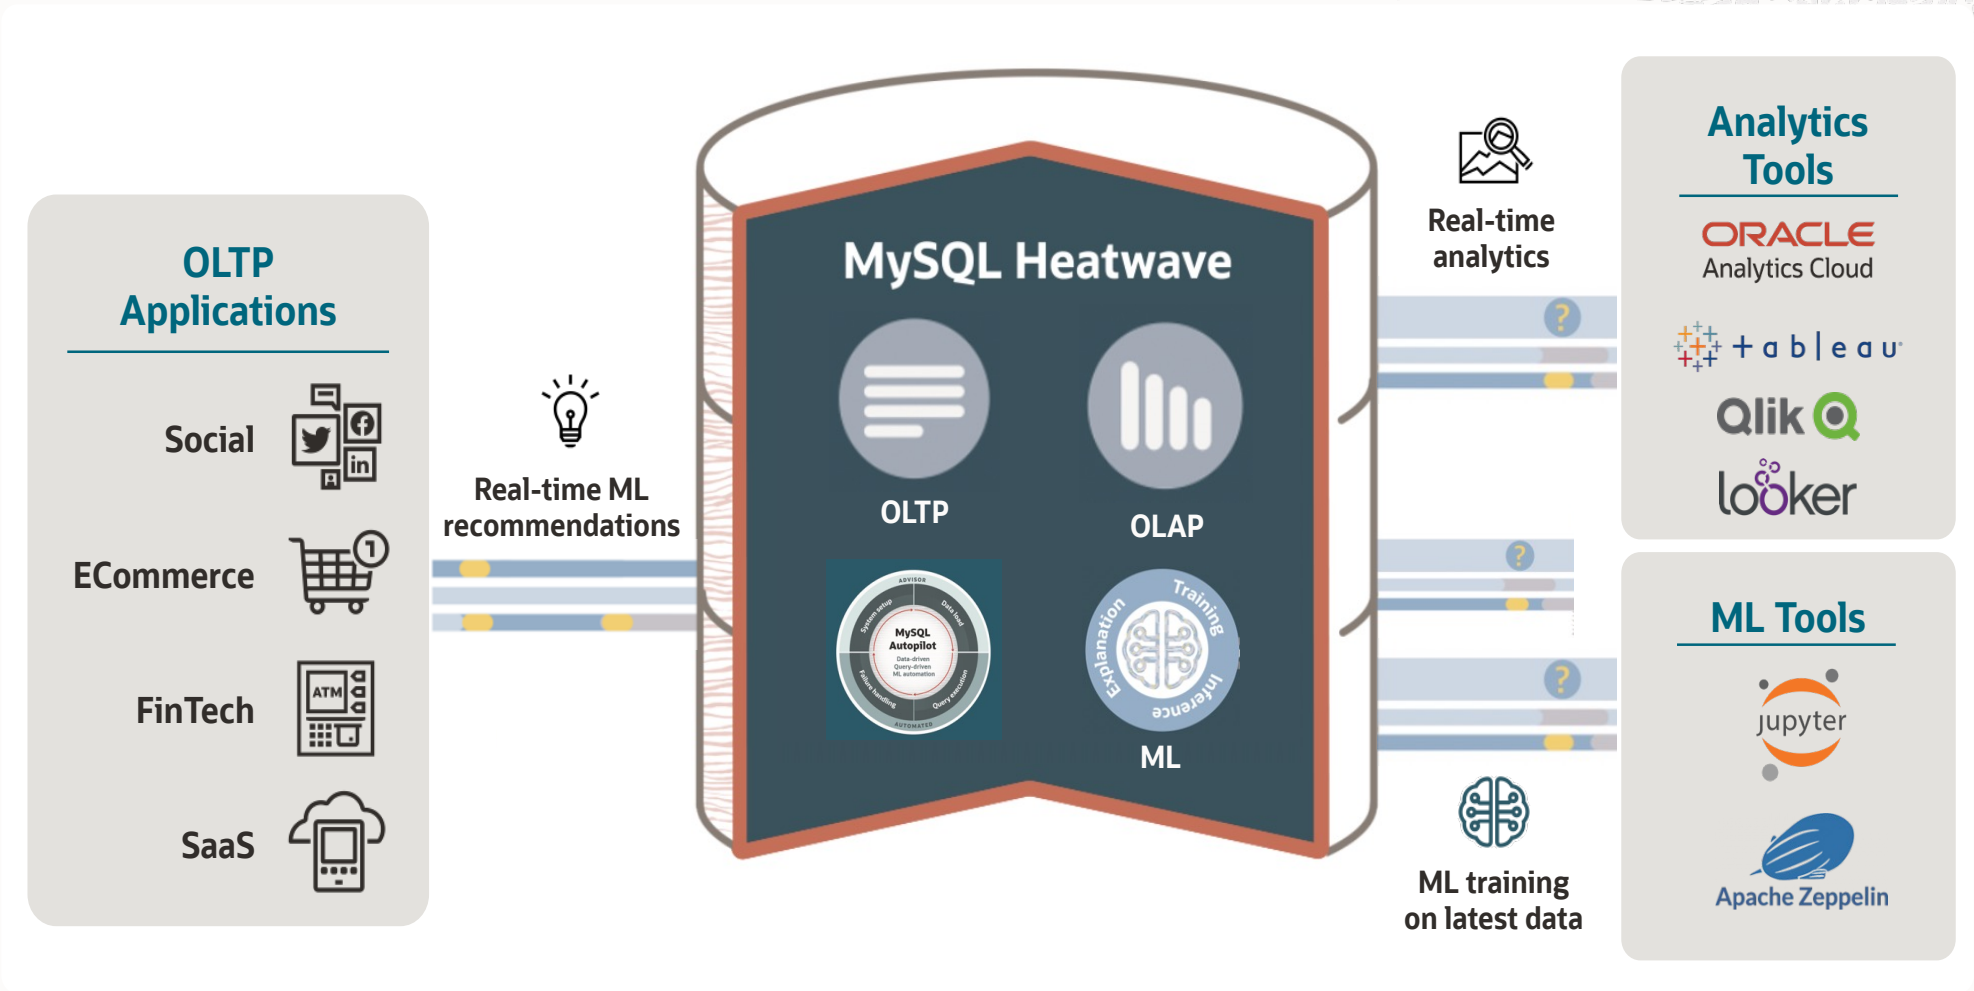
Online gaming

Users profile management

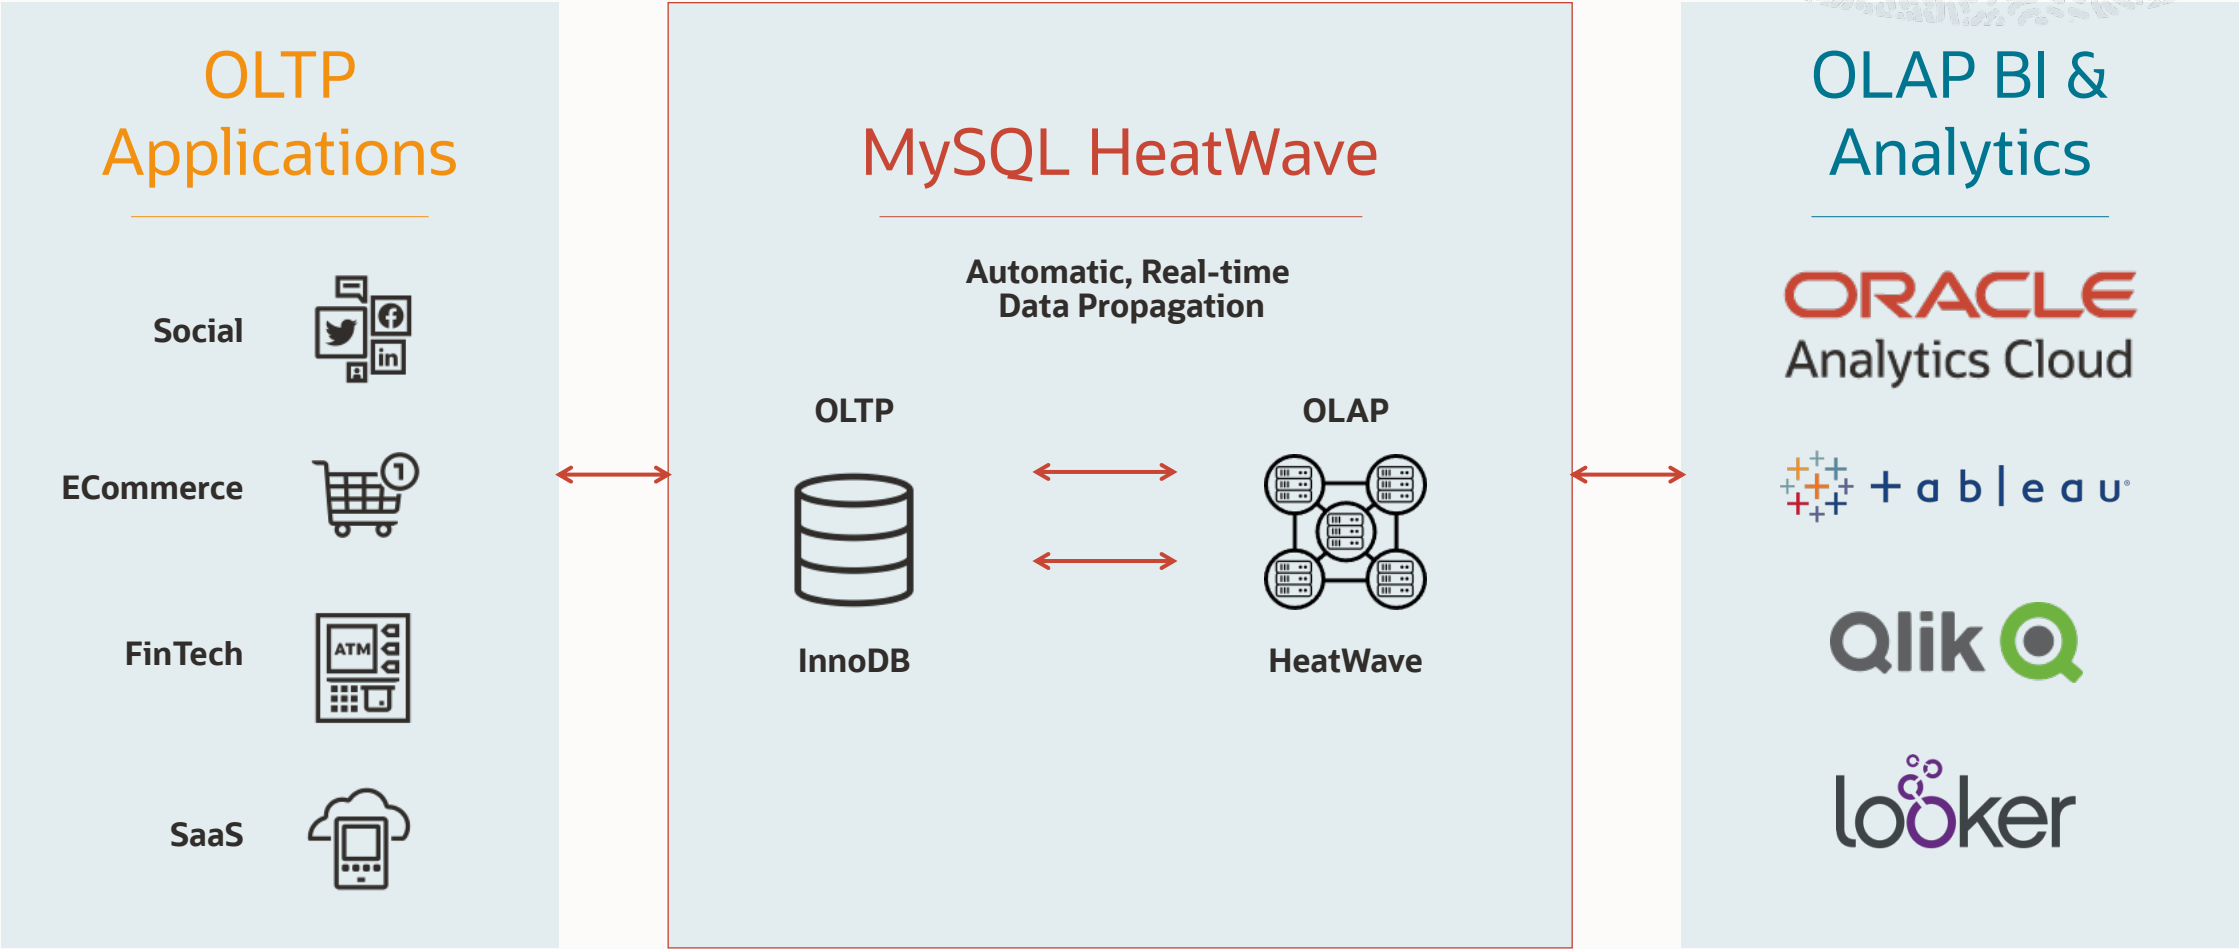
Employee portals

Logistics applications

MySQL HeatWave – One database for OLTP, OLAP, ML




MySQL compatible applications run without any changes



Price-Performance comparison for analytics

23x better than Redshift, 35x better than Snowflake, 41x better than Big Query, 15x better than Synapse

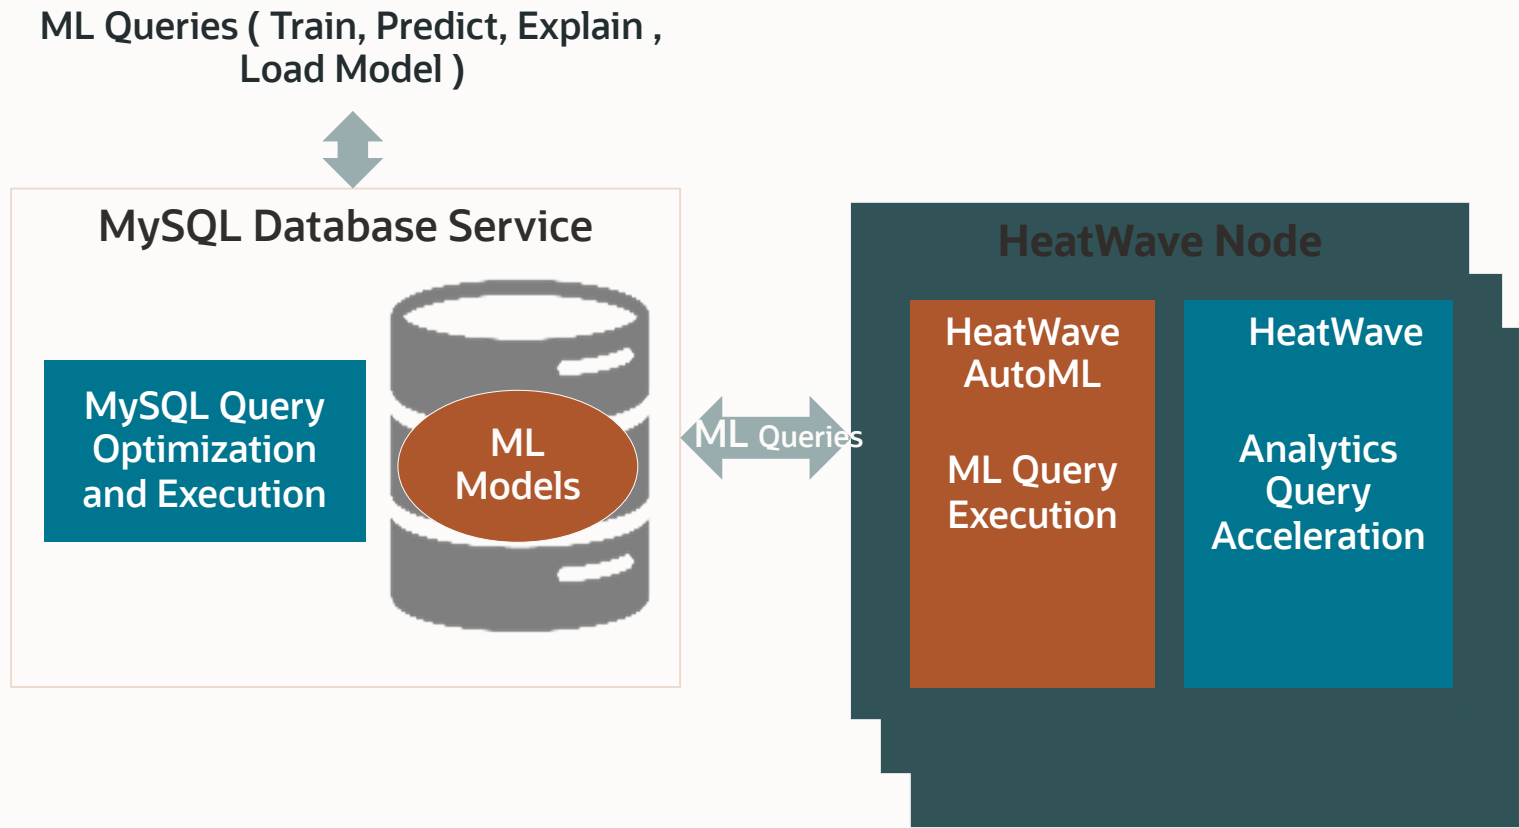


- Pricing for Redshift is based on 1-year reserved instance, paid upfront. For Snowflake is based on standard edition
- Pricing for Google Big Query is based on monthly flat rate commitment. For Azure Synapse is based on 1-year reserved pricing

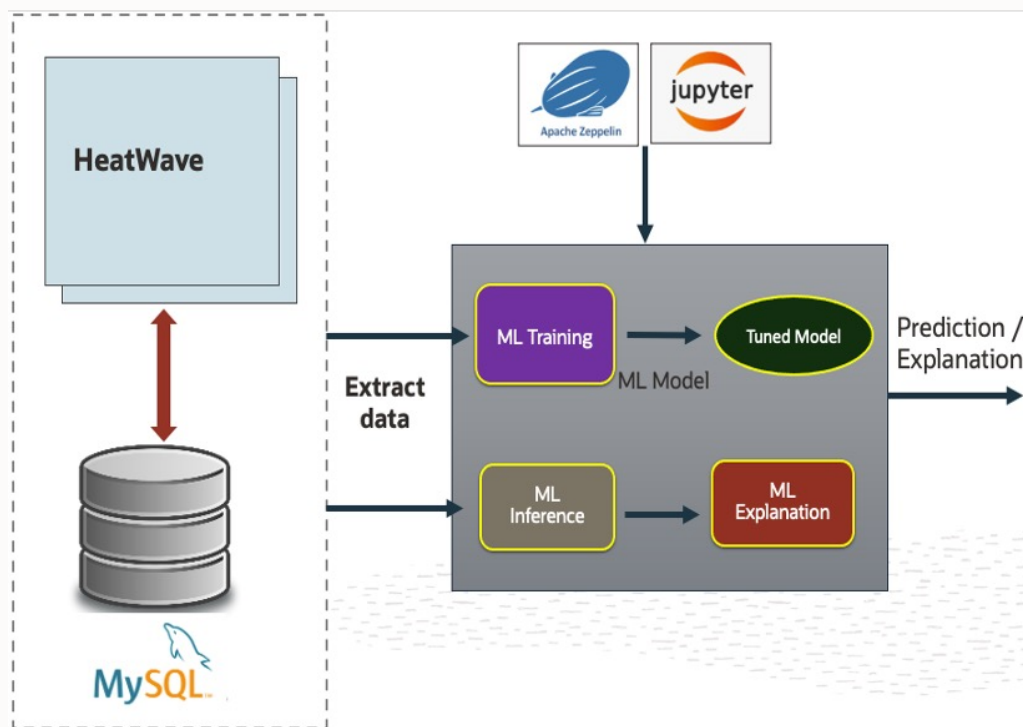
*Benchmark queries are derived from the TPC-H benchmarks, but results are not comparable to published TPC-H benchmark results since these do not comply with the TPC-H specifications.

HeatWave AutoML

- Fully Managed, highly scalable, cost-efficient machine learning solution
- No need to move data or ML models outside of the Database Service
- Database user friendly interfaces
- Leverages HeatWave cluster for distributed ML
- Model agnostic explanations for understanding both model and prediction behavior

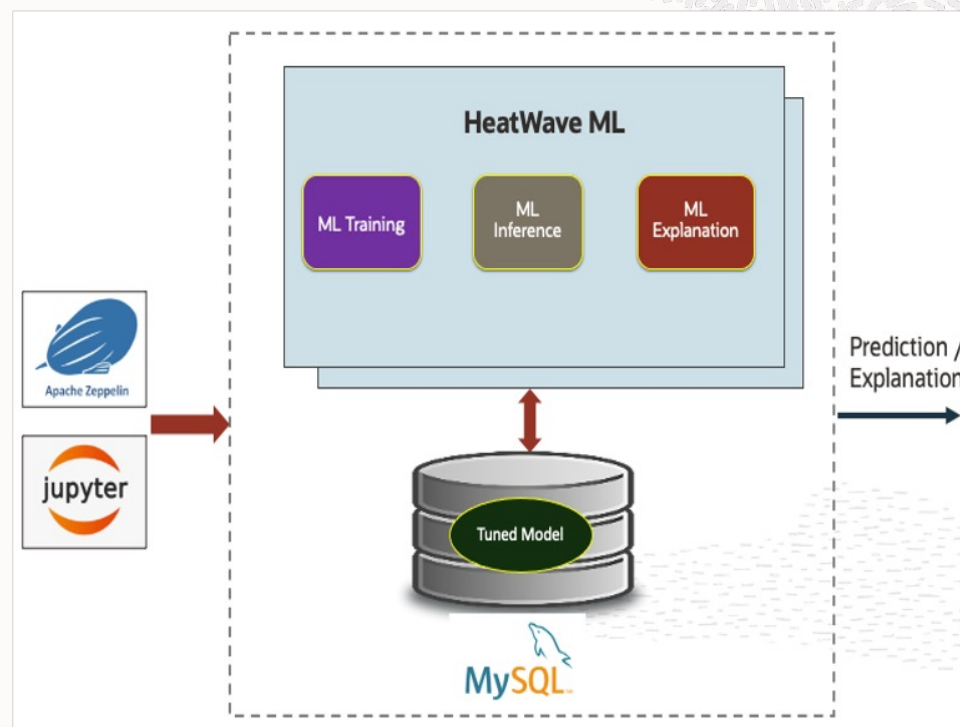


In-Database machine learning



Current ML processing approach

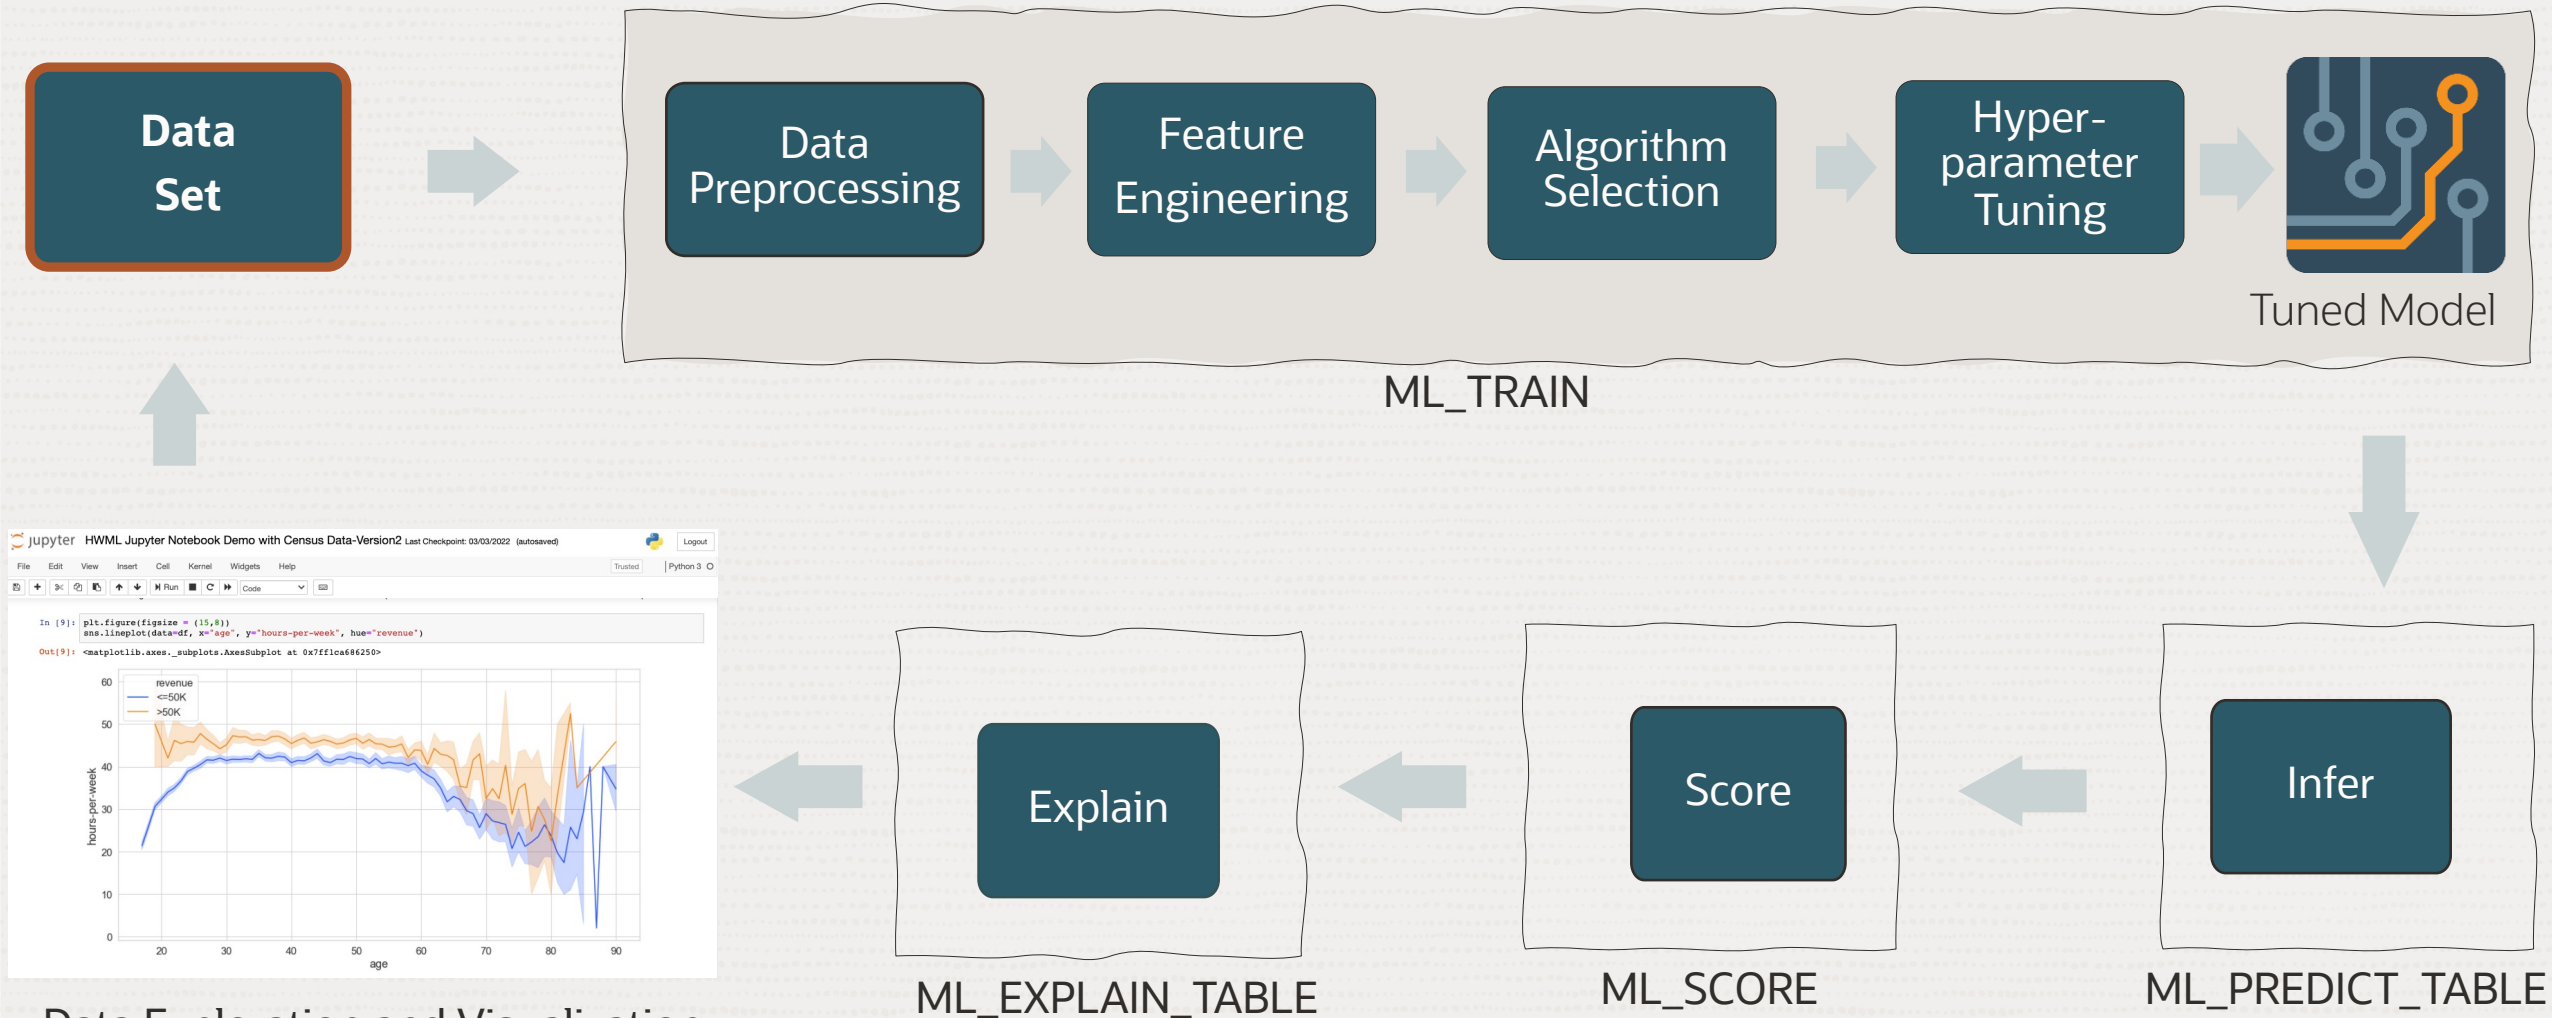
Requires ETL from database



Machine learning with HeatWave AutoML

No ETL, secure and faster

Model development with MySQL HeatWave AutoML



Data Exploration and Visualization

HeatWave ML vs Redshift ML: Benchmarks




- ✓ Produces **more accurate results**
- ✓ Trains models **25X faster** on average
- ✓ **1% of the cost**
- ✓ **Scales** as more modes are added

See Benchmark details: <https://www.oracle.com/mysql/heatwave/performance/>

Building MySQL Applications with Notebook

Why use notebook environment?




Community support

Adopted by data science community



Extensions and integrations

Supports multiple frameworks and languages




Free, Open source, Interactive


Can be used as IDE as well as presentation tool



Use popular notebooks to build your ML applications



Jupyter Notebook



The image shows the Apache Zeppelin Notebook interface. The title bar is blue with the Zeppelin logo and the text 'Notebook' and 'Job'. The main content area is titled 'Census-Demo' and contains a code cell with the following Python code:

```
%python
def show_model_explanations():
    import json
    df = execute_sql(f"""SELECT model_explanation FROM ML_SCHEMA_{DBSYSTEM_USER}.MOD
                        WHERE model_handle = @model""")
    return pd.Series(json.loads(df.iloc[0]['model_explanation'])).sort_values(0, asc

show_model_explanations()
```

Below the code, the output is displayed: 'Running SELECT model_explanation FROM ML_SCHEMA_salil.MODEL_CATALOG WHERE model_handle = @model' and '<AxesSubplot:>'. At the bottom, it says 'Took 1 sec. Last updated by anonymous at September 14 2022 1:16:11 PM'.

Apache Zeppelin Notebook



Connect to HeatWave Cluster

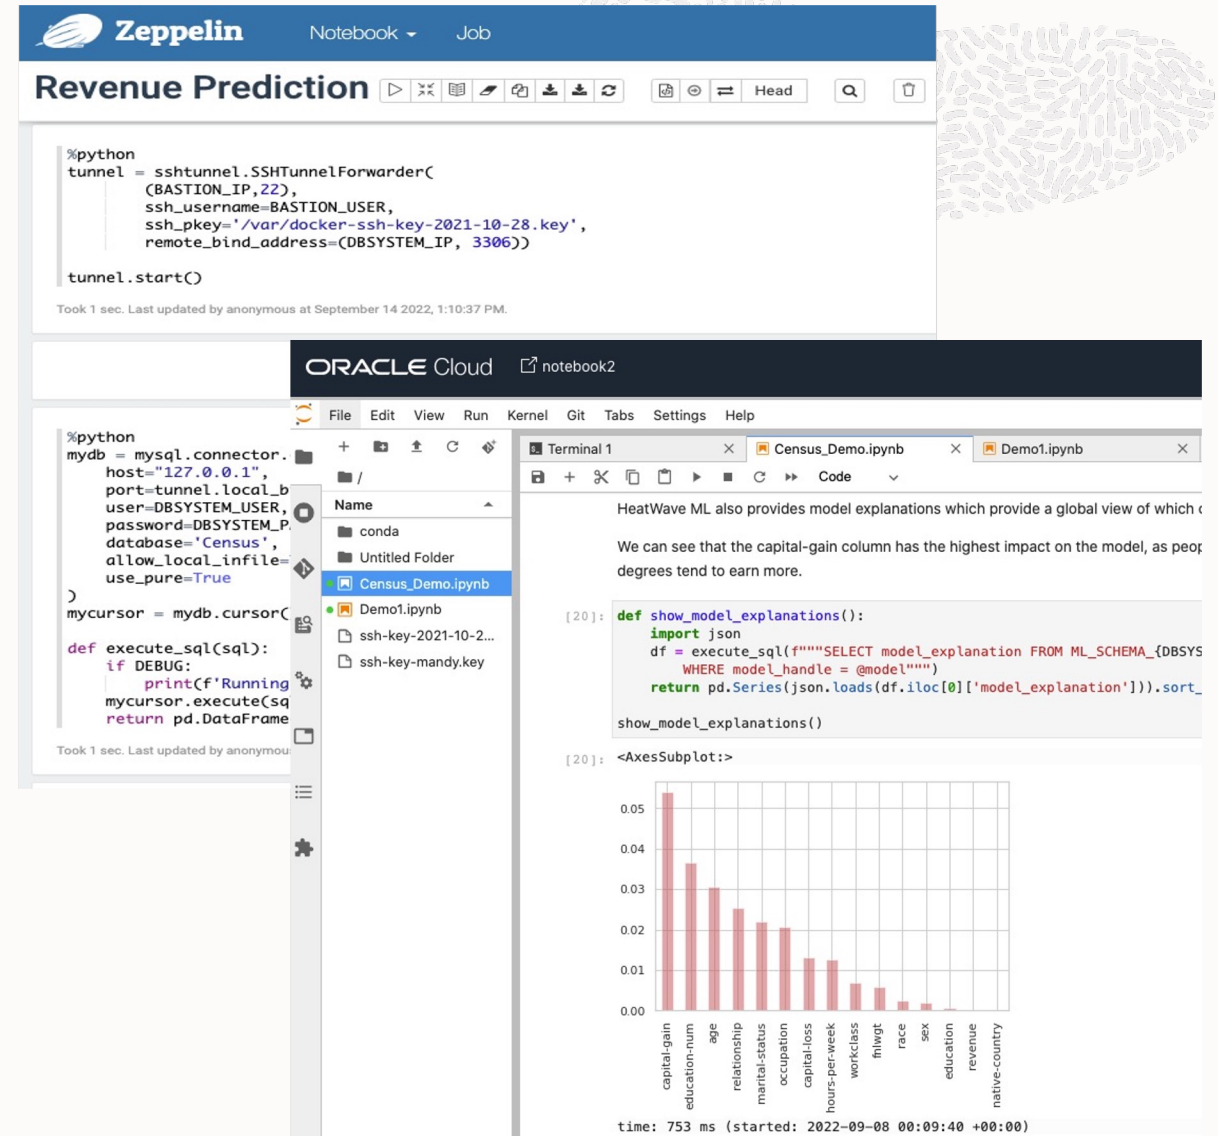
Create SSH Tunnel and establish connectivity with the HeatWave Cluster

Visualize and explore data

Identify critical data features, engineer new features

Invoke HeatWave MySQL HeatWave AutoML routines

Build, Score, Predict, Explain ML models and render the results



The screenshot displays the Zeppelin Notebook interface for a 'Revenue Prediction' job. The top navigation bar includes the Zeppelin logo, 'Notebook', and 'Job' tabs. The notebook content area shows Python code for setting up an SSH tunnel and connecting to a MySQL database. The code is as follows:

```
%python
tunnel = sshtunnel.SSHTunnelForwarder(
    (BASTION_IP, 22),
    ssh_username=BASTION_USER,
    ssh_pkey='/var/docker-ssh-key-2021-10-28.key',
    remote_bind_address=(DBSYSTEM_IP, 3306))

tunnel.start()
```

Below the code, it indicates 'Took 1 sec. Last updated by anonymous at September 14 2022, 1:10:37 PM.' The interface also features a file explorer on the left, a terminal window, and a bar chart visualization of model explanations. The bar chart shows the impact of various features on the model, with 'capital-gain' having the highest impact.

HeatWave ML also provides model explanations which provide a global view of which features have the most impact on the model. We can see that the capital-gain column has the highest impact on the model, as people with higher capital gain tend to earn more.

```
[20]: def show_model_explanations():
import json
df = execute_sql(f"""SELECT model_explanation FROM ML_SCHEMA.{DBSYS
WHERE model_handle = @model""")
return pd.Series(json.loads(df.iloc[0]['model_explanation'])).sort_

show_model_explanations()
```

[20]: <AxesSubplot:>

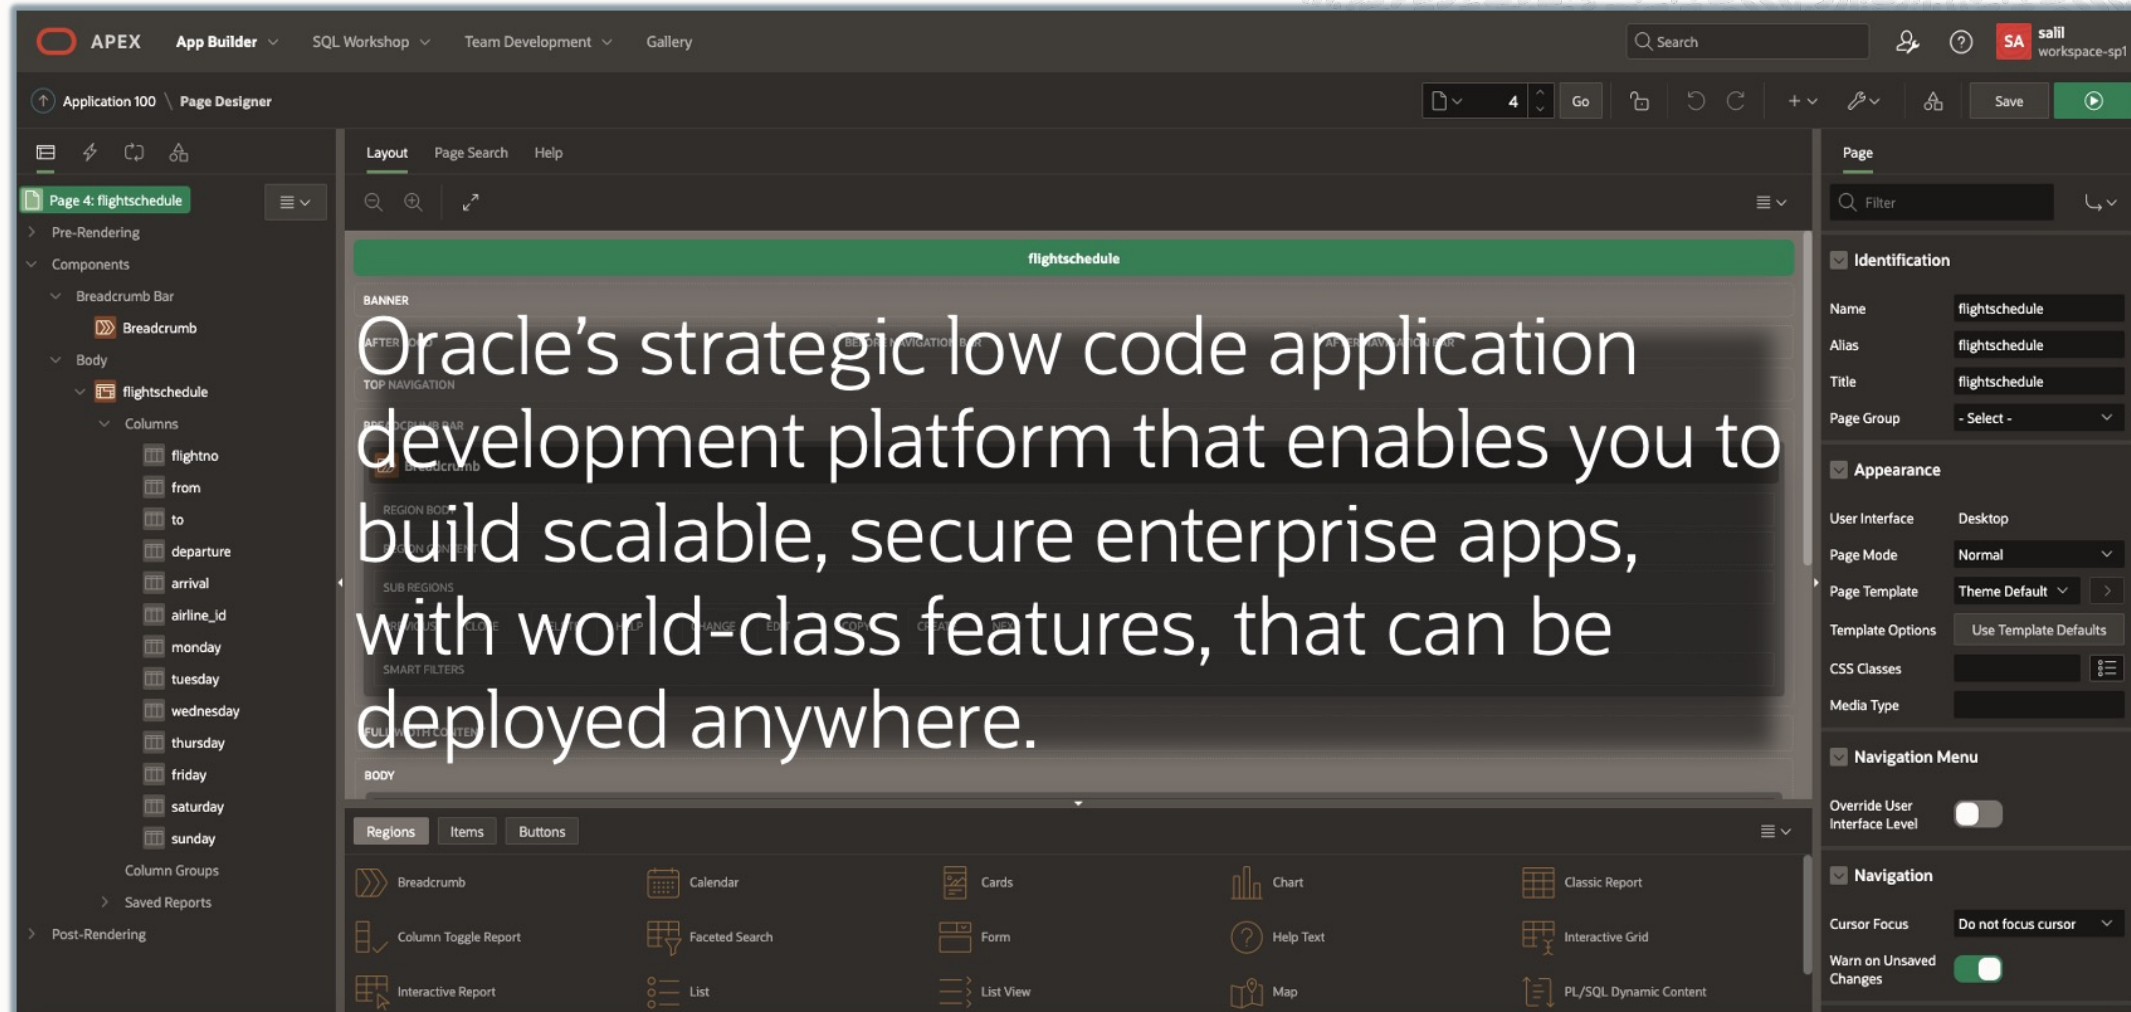
time: 753 ms (started: 2022-09-08 00:09:40 +00:00)



Product Demo

Building MySQL Applications with APEX

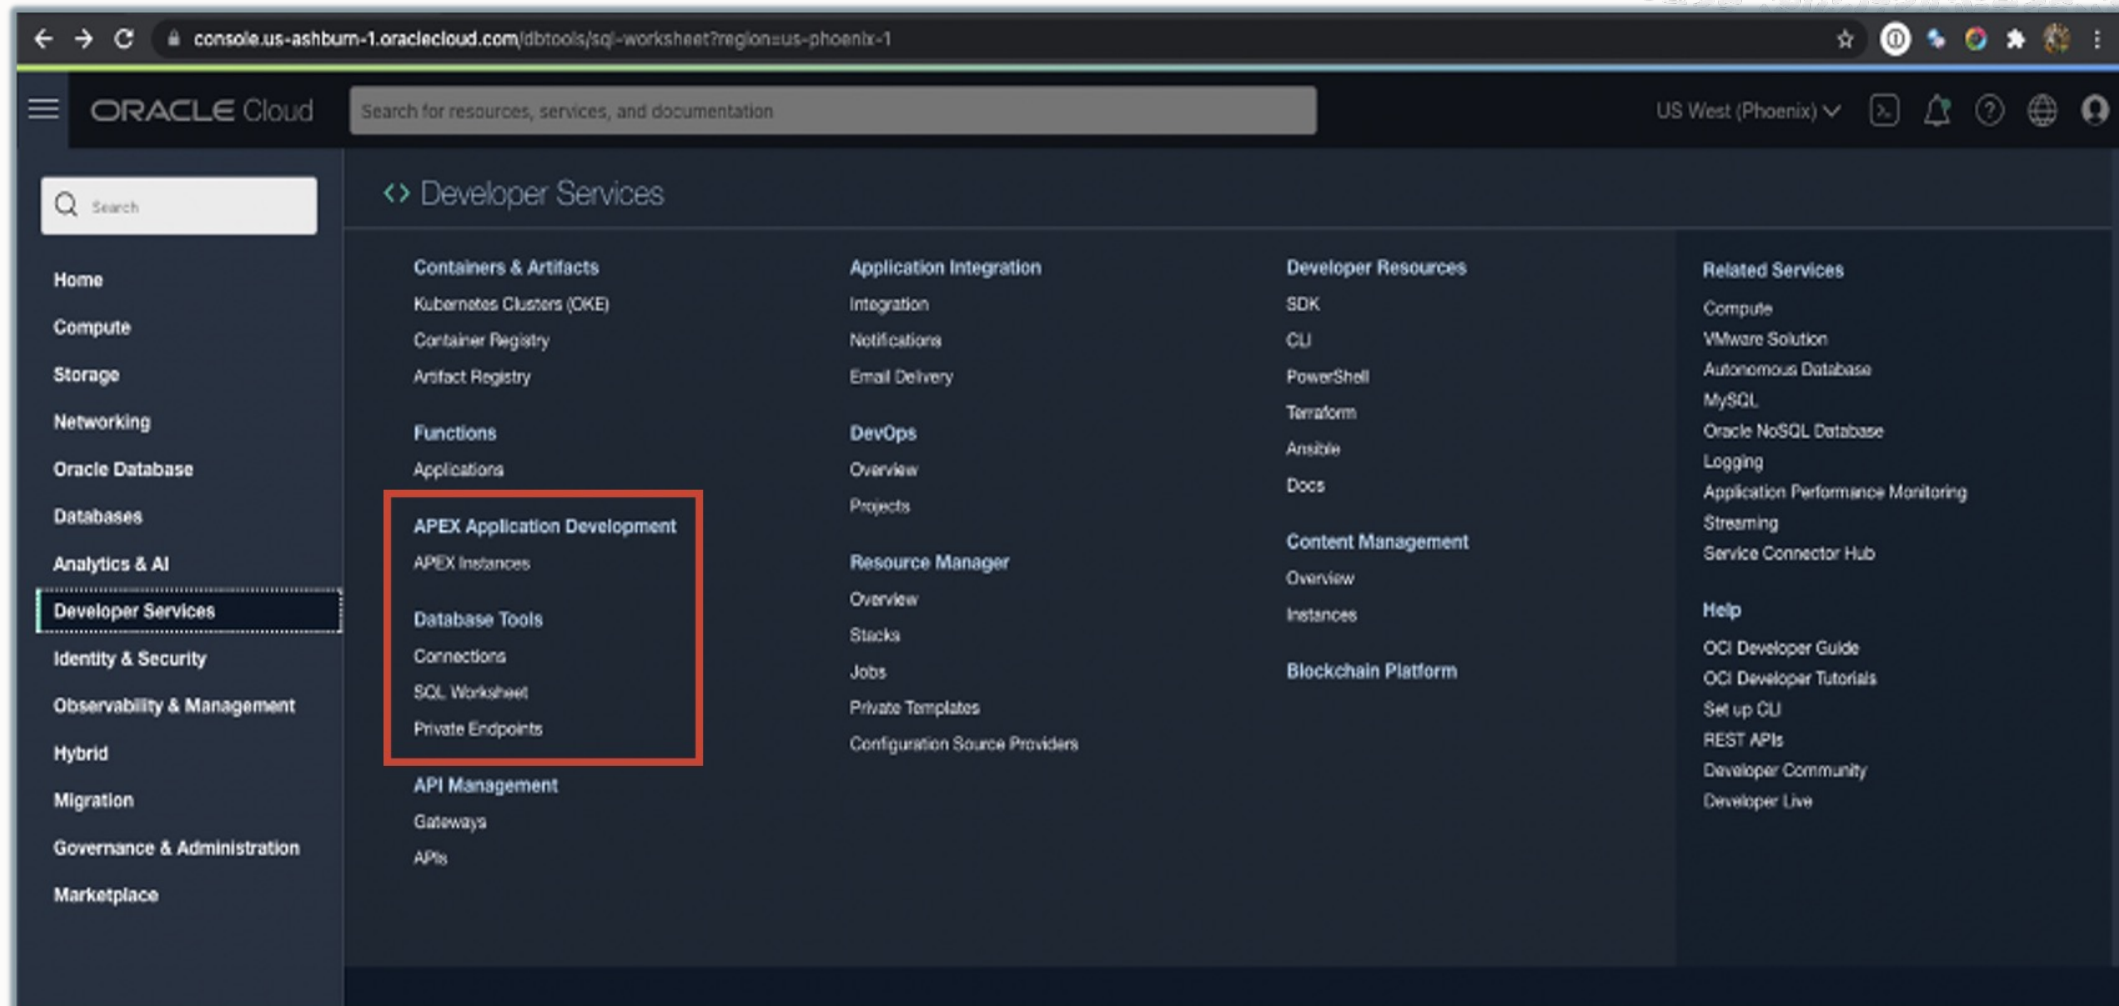
What is APEX?




The screenshot displays the Oracle APEX App Builder interface. The top navigation bar includes the APEX logo, 'App Builder', 'SQL Workshop', 'Team Development', and 'Gallery'. A search bar and user profile 'SA salil workspace-sp1' are on the right. The main workspace shows 'Page 4: flightschedule' in the 'Page Designer' view. The left sidebar lists components like 'Pre-Rendering', 'Components', 'Breadcrumb Bar', 'Body', 'Columns', and 'Post-Rendering'. The central canvas shows a flight schedule layout with sections like BANNER, TOP NAVIGATION, REGION BODY, SUB REGIONS, SMART FILTERS, and BODY. The right sidebar contains settings for 'Page', 'Identification', 'Appearance', 'Navigation Menu', and 'Navigation'. The bottom of the interface features a 'Regions' tab and a palette of components including Breadcrumb, Calendar, Cards, Chart, Classic Report, Column Toggle Report, Faceted Search, Form, Help Text, Interactive Grid, Interactive Report, List, List View, Map, and PL/SQL Dynamic Content.

Oracle's strategic low code application development platform that enables you to build scalable, secure enterprise apps, with world-class features, that can be deployed anywhere.

Database tools for connect management and SQL worksheet




APEX distinguishing characteristics



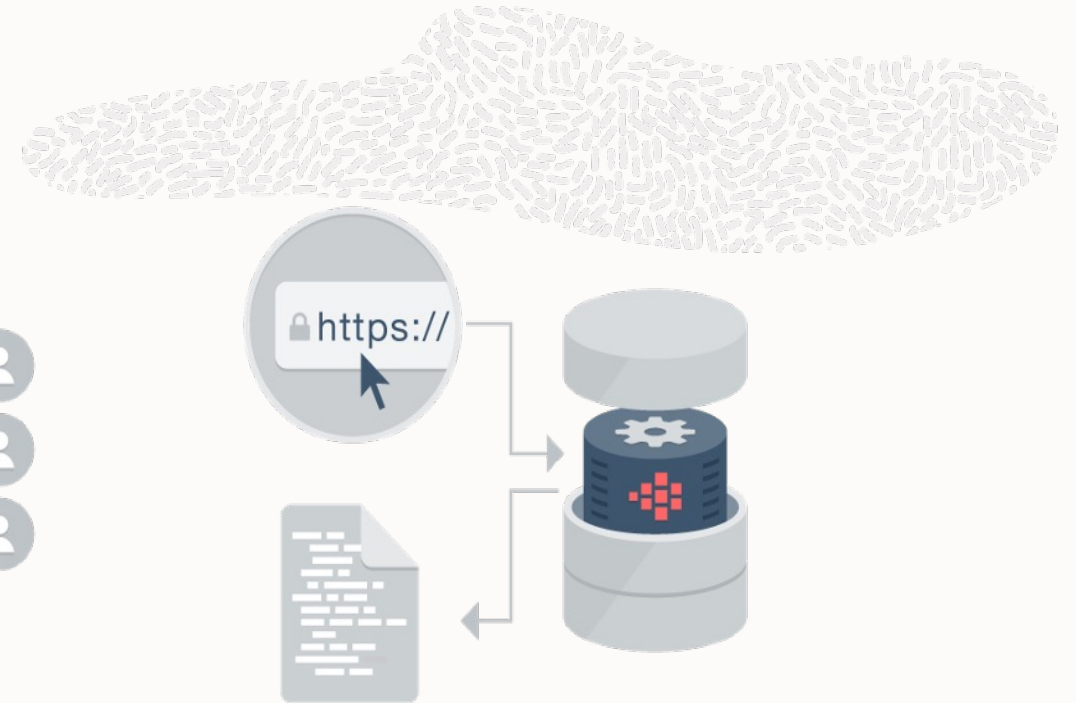
App Development IDE is a web browser.

No client software needed



App definitions are stored in the database as meta data.

Declarative – No code generation




Page generation is efficient with only one request and one response.

Data processing done in the Database

APEX integration with MySQL Database Service

- Based on REST Enabled SQL provided by ORDS
- APEX talks to ORDS (Send SQL over HTTP / REST)
- ORDS talks to remote MySQL database on Oracle Cloud Infrastructure





Product Demo

Building MySQL Applications with Oracle Analytics Cloud

Oracle Analytics Cloud



**Data
Engineers**



**Citizen
Data Scientists**



**Business
Analysts**

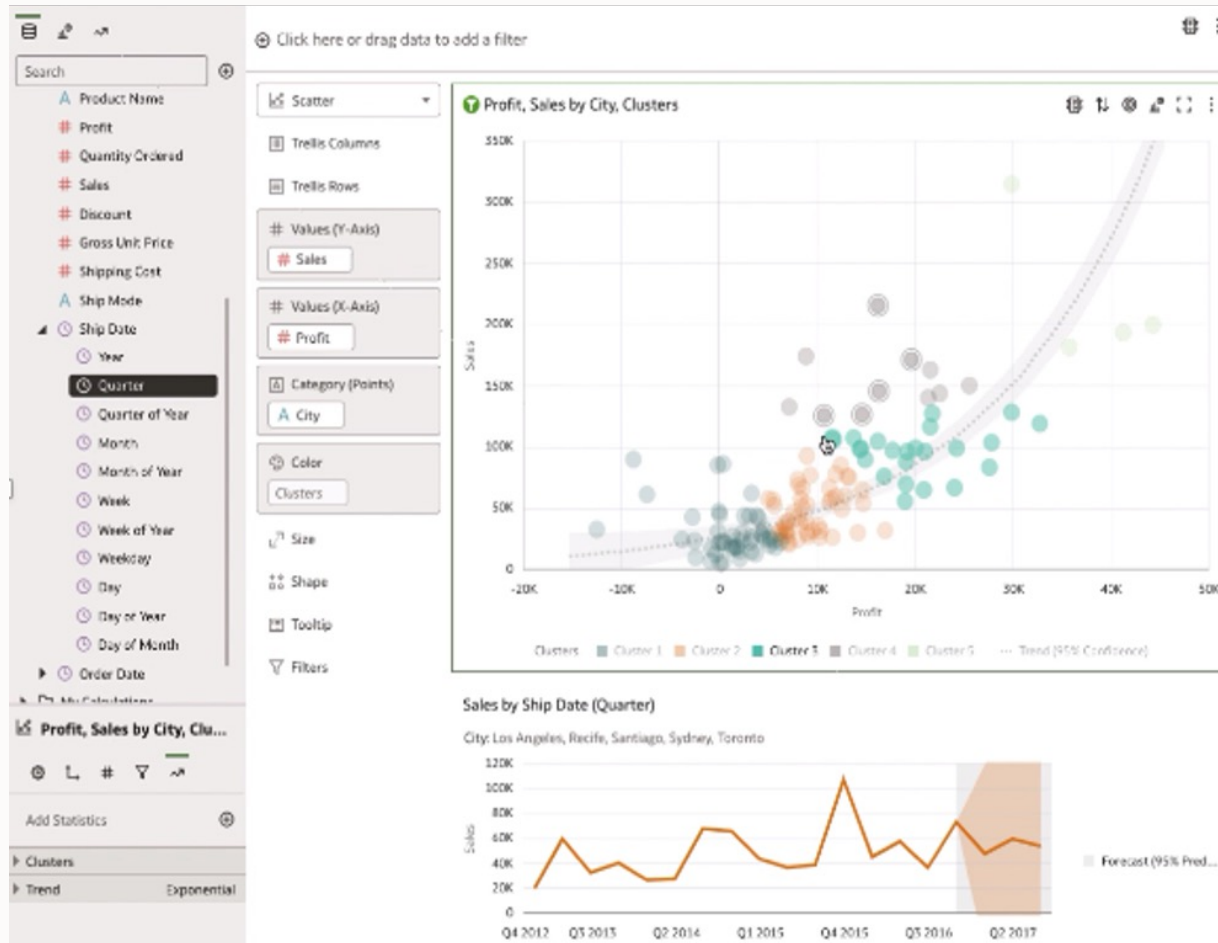


**Business
Users**

Governed Analytics		Self-Service Analytics		Augmented Analytics	
Dashboards	Distributed Reports	Data Preparation	Data Visualization	Voice & Chatbot	Natural Language
Semantic Models	Query Federation	Storytelling	Collaboration	Data Enrichment	1-Click Explain
Role Based Security	Data Export	What-If Analysis	Mobile	Adaptive & Personalized	Machine Learning
Open Connectivity Oracle Cloud Third Party Cloud On-Premises					

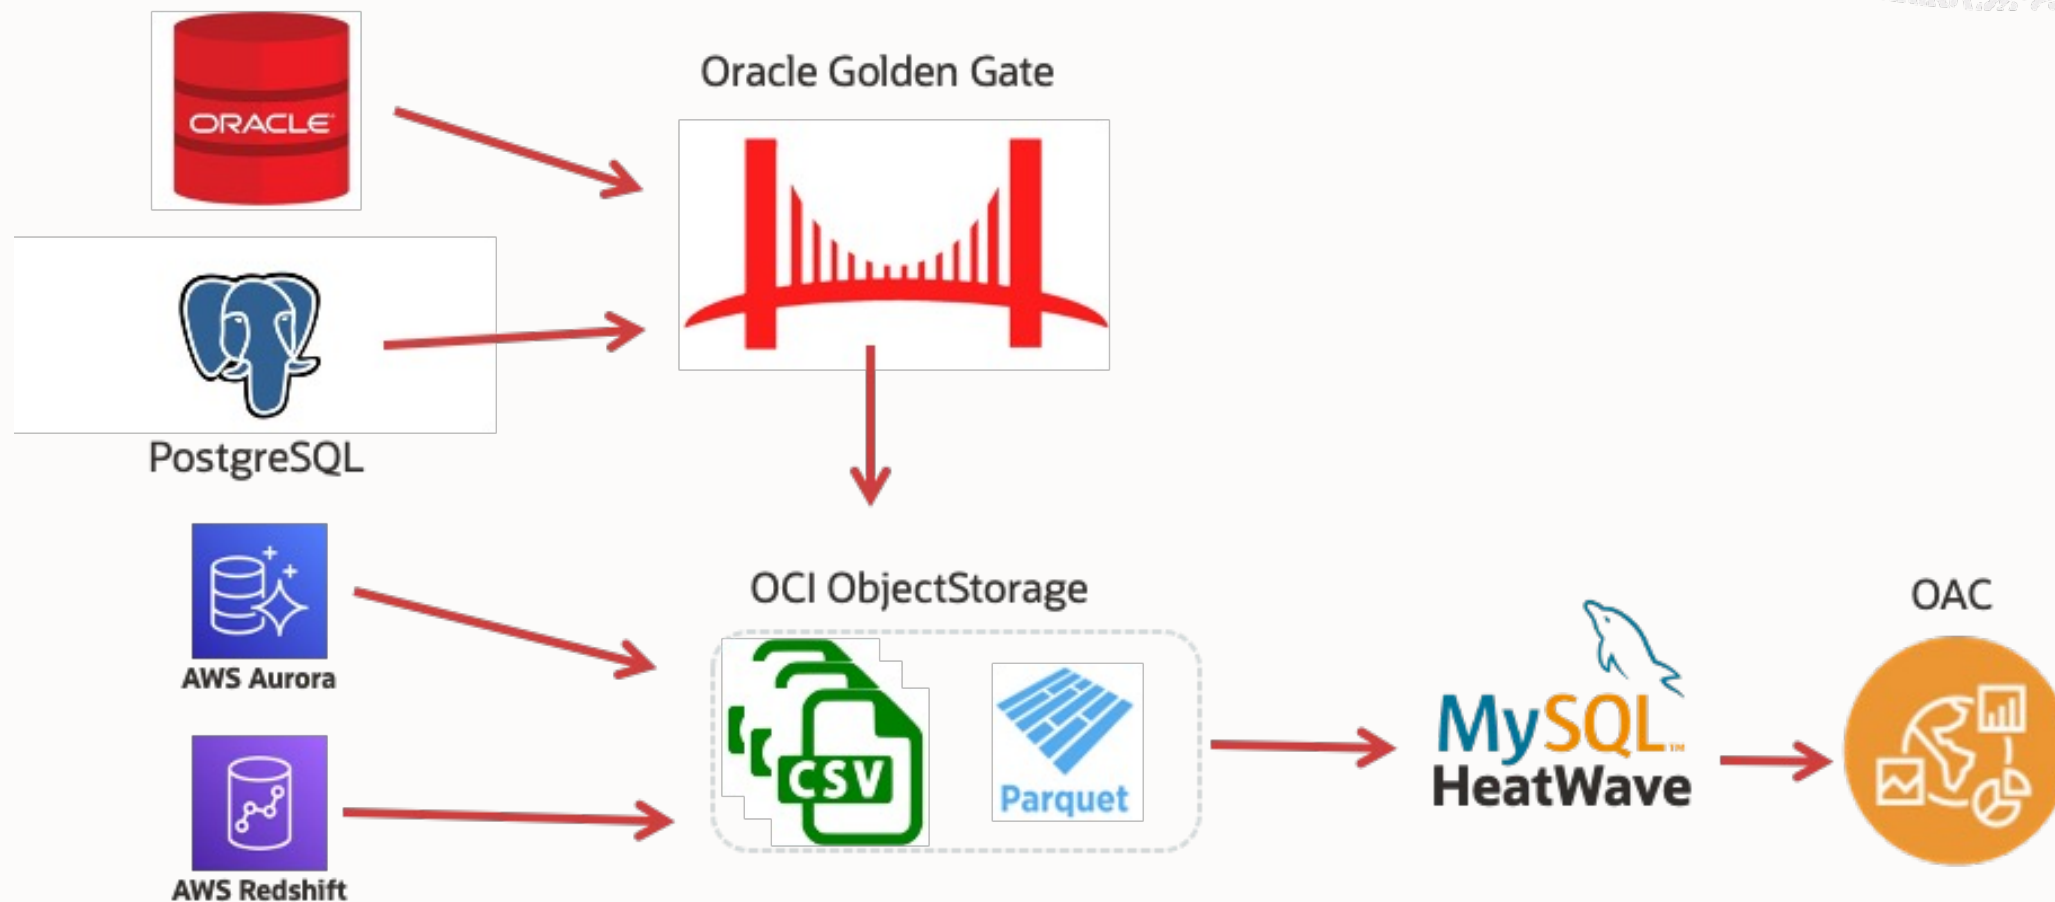


Oracle Analytics Cloud with MySQL Database Service



- Optimized native connection to MySQL
- High performance interactive analytics
- Supports connectivity to private MySQL Service
- Allow MySQL users to take full advantage of all capabilities that OAC offers

MySQL HeatWave can process data from multiple data sources



MySQL REST Service (MRS)

Fast, Secure HTTPS Access for your MySQL Data



- Enables fast and secure HTTPS access for your MySQL data
- Implemented as a MySQL Router feature
- Publish RESTful Web Services for interacting with the data stored in MySQL solutions

RESTful Web Services

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

MySQL Shell for VS Code

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

Built in User Management

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

MySQL REST Service: Architecture

Fast, Secure HTTPS Access for your MySQL Data

