



Register for the Select AI Office Hours series

Select AI Office Hours

Getting Started with Select AI

Mark Hornick
Product Management
Senior Director
Machine Learning and AI

Sherry LaMonica
Product Management
Consulting MTS

February 7, 2025



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

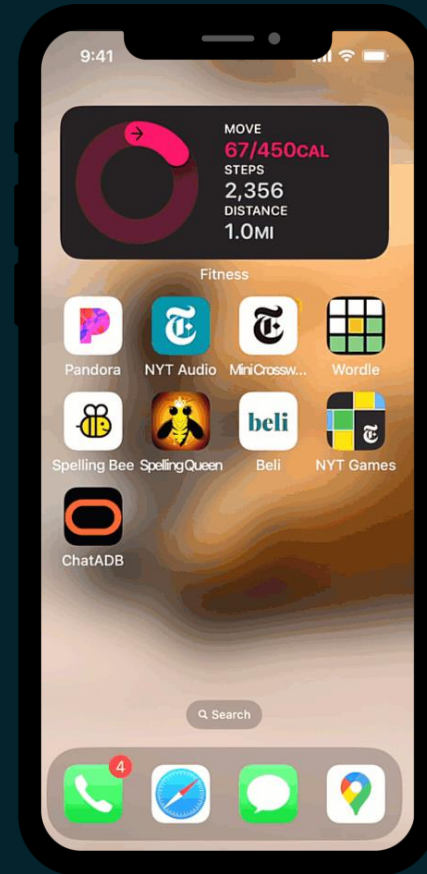
The materials in this presentation pertain to Oracle Health, Oracle, Oracle Cerner, and Cerner Enviza which are all wholly owned subsidiaries of Oracle Corporation. Nothing in this presentation should be taken as indicating that any decisions regarding the integration of any EMEA Cerner and/or Enviza entities have been made where an integration has not already occurred.

Agenda

- Select AI introduction
- NL2SQL
- Chat genAI use case
- RAG
- Synthetic Data Generation
- New features
- Roadmap

Select AI

Supporting LLM-enhanced application development and productivity through natural language

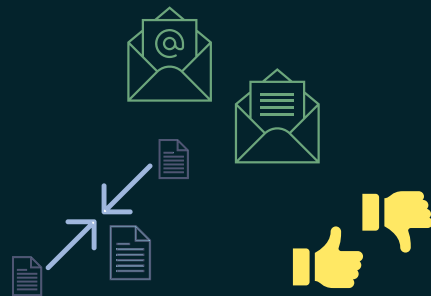


What can you do with **Select AI**?

A SQL interface for leveraging generative AI on Oracle Autonomous Database

Create applications using generative AI

- customer email generation
- recommendations
- sentiment analysis
- document summarization
- ...and more



Generate SQL specific to your database as a developer assistant

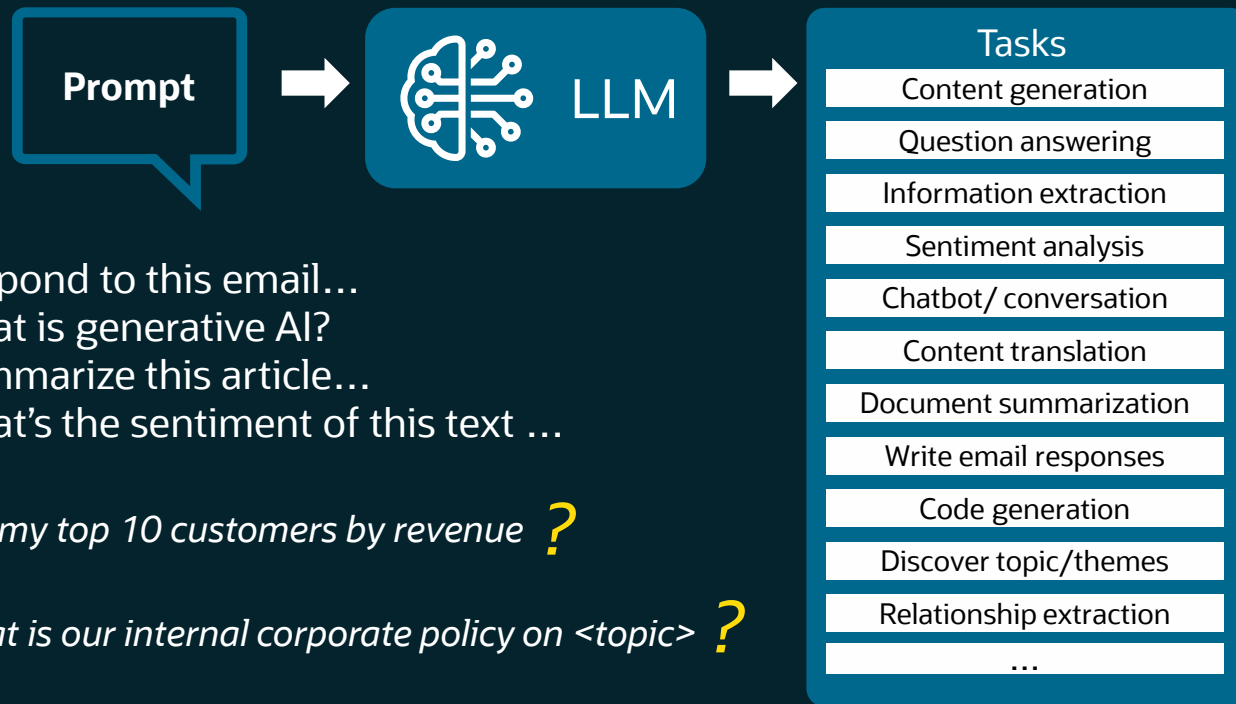
Build apps to interact with data using natural language

Generate RAG-enhanced results using vector search

Generate synthetic data

Large Language Models

Generate text content – “next word predictors”



- Respond to this email...
- What is generative AI?
- Summarize this article...
- What's the sentiment of this text ...
- *List my top 10 customers by revenue ?*
- *What is our internal corporate policy on <topic> ?*

Issues with LLMs

Hallucination

A confident response that doesn't seem justified by its training data or is simply not true

Bots can embed plausible-sounding random false information in their generated content

Often happens when an LLM doesn't understand the context of the request

Bias

Can reflect or augment biases present in training data text

E.g., if the training data contains articles on physicians, and the text mentions only doctors of a single gender, the LLM's responses will reflect this

Limitations

Can only know information available up to the training date of the LLM

General knowledge – not about your business/enterprise

Not “aware” of you as an individual or your personal data

Tokens limits on input prompt and output can impact session “memory” when generating responses

Select AI can help address such issues ...



ChatDB Demonstration

Using an APEX-based chatbot for NL2SQL

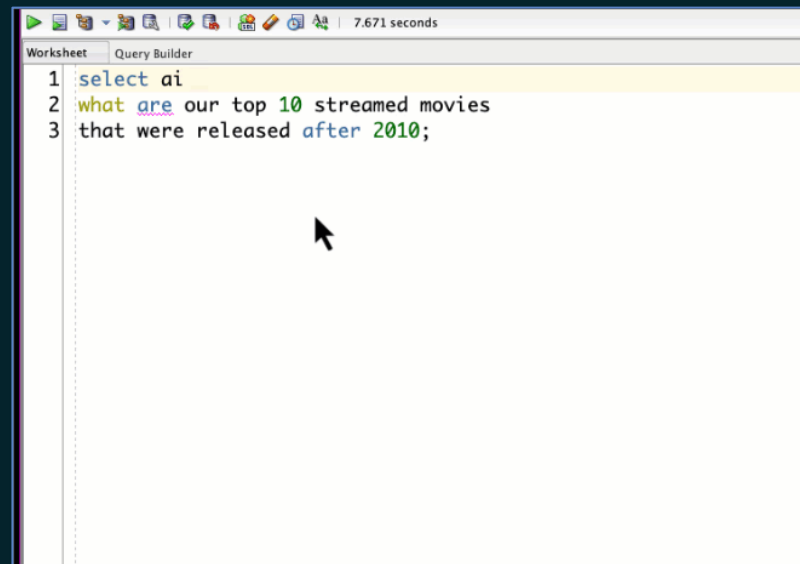
Use natural language to query data and get responses using generative AI

Easily access LLMs from multiple providers

- OCI GenAI Service
- OpenAI
- Azure OpenAI Service
- Cohere
- Google
- Anthropic
- Hugging Face
- ...more coming

Actions

- `runsql` - return the SQL result set (default)
- `showsql` - return the generated query
- `explainsql` - explain generated SQL query
- `showprompt` - display the generated prompt
- `narrate` - return a conversational result
- `chat` - general AI chat - passthrough to LLM



The screenshot shows a 'Query Builder' window with a 'Worksheet' tab. The interface displays a natural language query in three lines, which is being converted into a SQL query. The natural language query is: 'what are our top 10 streamed movies that were released after 2010;'. The corresponding SQL query is: 'select ai'. The window title bar shows 'Worksheet Query Builder' and a timer indicating '7.671 seconds'.

```
1 select ai
2 what are our top 10 streamed movies
3 that were released after 2010;
```

SQL generation

ChatGPT Example

Joins invented
"Movies" and
"StreamData"
tables

Aggregates, including
movie name

Orders result and
selects top 10

Generate the SQL to find our top 10 streamed movies

To find the top 10 streamed movies, you'd likely have a table that tracks streaming data. Let's assume you have a table called `StreamData` with columns `MovieID` and `StreamCount` (indicating the number of times each movie has been streamed).

Another table called `Movies` might contain information about each movie, including a unique `MovieID` and `MovieName`.

Here is an SQL query that will find the top 10 streamed movies:

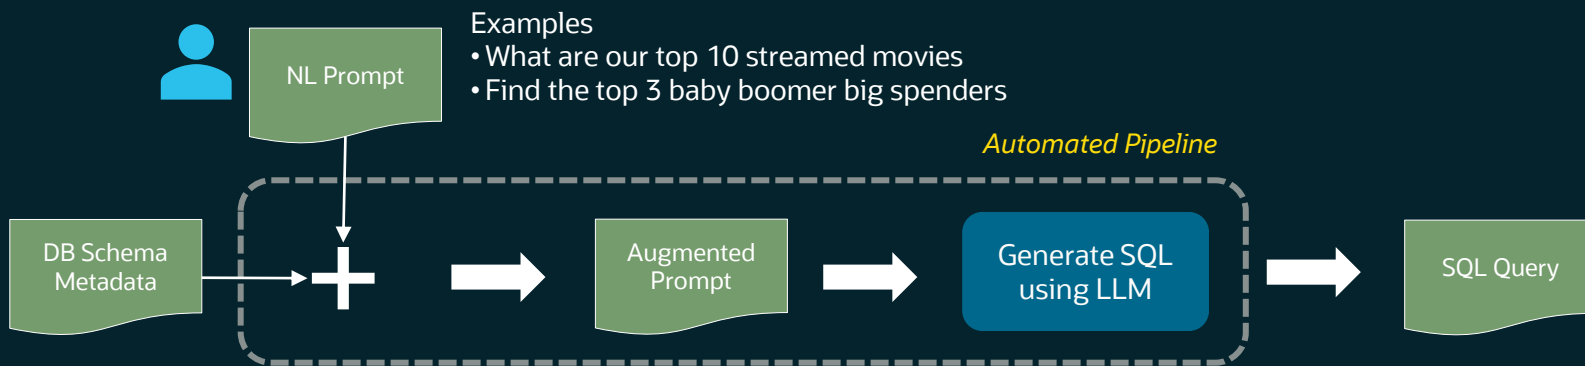
```
sql Copy code
SELECT m.MovieID, m.MovieName, SUM(s.StreamCount) AS TotalStreams
FROM Movies m
JOIN StreamData s ON m.MovieID = s.MovieID
GROUP BY m.MovieID, m.MovieName
ORDER BY TotalStreams DESC
LIMIT 10;
```

Uses "sum" for total streams

In this SQL query, the `Movies` table is joined with the `StreamData` table on the `MovieID` field. It groups the results by `MovieID` and `MovieName` to calculate the sum of `StreamCount` for each movie. The query then sorts these results in descending order based on `TotalStreams` and limits the output to 10 records to get the top 10 streamed movies.

Again, please note that this is based on a hypothetical database schema, so your actual query might need to be different based on your specific tables and columns.

SQL query generation pipeline with 'showsqli' action



Examples

- What are our top 10 streamed movies
- Find the top 3 baby boomer big spenders

Metadata sent to LLM

- Table names
- Column names and data types
- Comments on tables and columns



```
SQL> SELECT AI SHOWSQL what are our top 10 streamed movies
```

Query data in your database using natural language

An example for movie streaming service



Movies

Actors

Streams

Genre

Customer_extension

Customer_segment

Pizza_shop

Customer_contact

AI Profile

Configure how you want to Select AI to behave for SQL generation

Choose your AI provider

Choose your LLM

Work with your full schema or specific tables and views

Create profile →

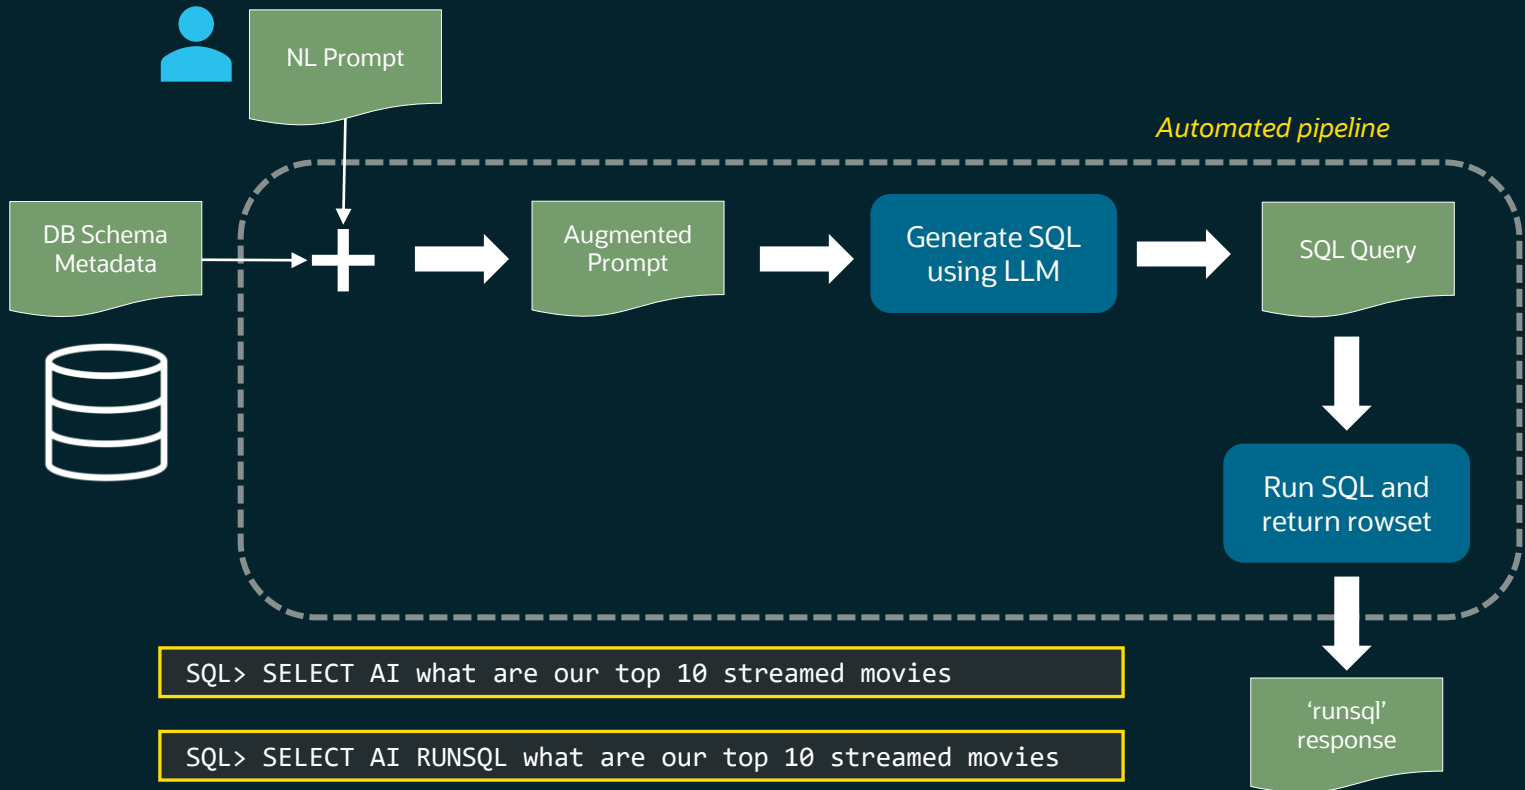
- Profile name
- AI provider
- Credential
- Object list

Set profile →

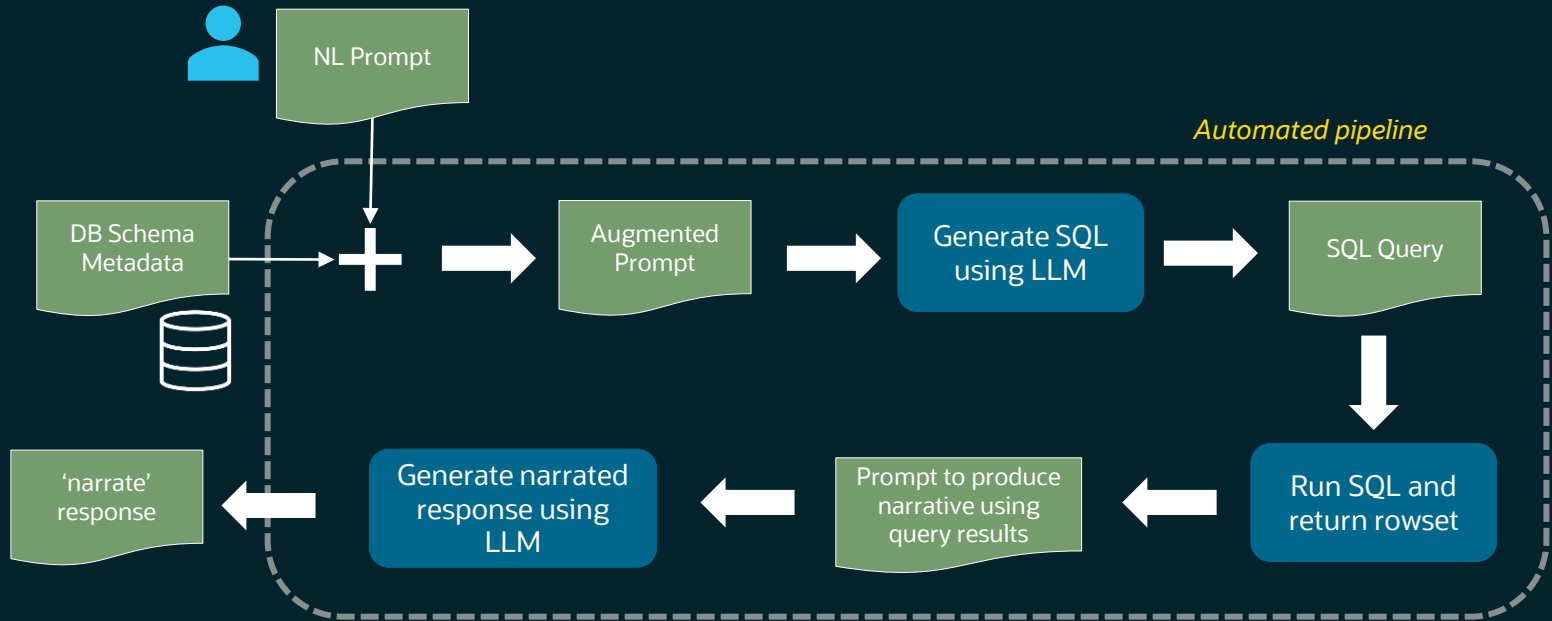
```
BEGIN
  DMBS_CLOUD_AI.create_profile(
    profile_name => 'openai_gpt4',
    attributes =>
      '{"provider": "openai",
       "credential_name": "OPENAI_CRED",
       "model": "gpt-4o",
       "object_list": [{"owner": "MOVIESTREAM", "name": "movies"},
                       {"owner": "MOVIESTREAM", "name": "streams"},
                       {"owner": "MOVIESTREAM", "name": "actors"},
                       {"owner": "MOVIESTREAM", "name": "genre"},
                       {"owner": "MOVIESTREAM", "name": "customer_segment"}]}'
  );

  dbms_cloud_ai.set_profile(profile_name => 'openai_gpt4');
END;
```

SQL query generation pipeline with 'runsql' action

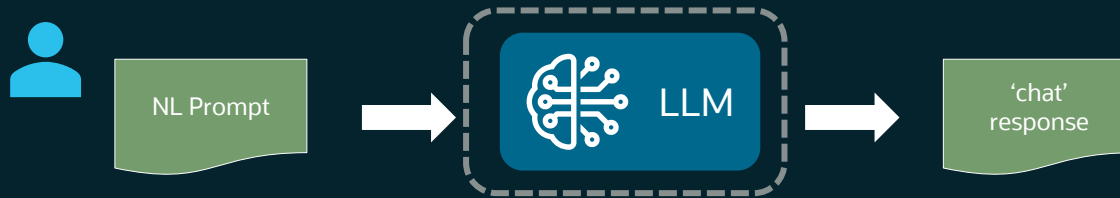


SQL query generation pipeline with 'narrate' action



```
SQL> SELECT AI NARRATE what are our top 10 streamed movies
```

Select AI pipeline with 'chat' action



```
SQL> SELECT AI CHAT what is Oracle Autonomous Database
```




Select AI Demonstration

Getting started and using NL2SQL

Better SQL generation through better metadata

Metadata is key – a pathological case

TABLE1	TABLE2	TABLE3
<pre>CREATE TABLE table1 (c1 NUMBER, c2 VARCHAR2(200), c3 NUMBER)</pre>	<pre>CREATE TABLE table2 (c1 TIMESTAMP, c2 NUMBER, c3 NUMBER, c4 NUMBER, c5 VARCHAR2(100), c6 NUMBER, c7 NUMBER)</pre>	<pre>CREATE TABLE table3 (c1 NUMBER, c2 VARCHAR2(30))</pre>

TABLE1
<pre>COMMENT ON TABLE table1 IS 'Contains movies, movie titles and the year it was released'; COMMENT ON COLUMN table1.c1 IS 'movie ids. Use this column to join to other tables'; COMMENT ON COLUMN table1.c2 IS 'movie titles'; COMMENT ON COLUMN table1.c3 IS 'year the movie was released';</pre>
TABLE2
<pre>COMMENT ON TABLE table2 IS 'transactions for movie views - also known as streams'; COMMENT ON COLUMN table2.c1 IS 'day the movie was streamed'; COMMENT ON COLUMN table2.c2 IS 'movie ids. Use this column to join to other tables';</pre>

```
begin  
  dbms_cloud_ai.create_profile(  
    profile_name => 'myprofile',  
    attributes =>  
      '{"provider": "azure",  
        "azure_resource_name": "my_resource",  
        "azure_deployment_name": "my_deployment",  
        "credential_name": "my_credential",  
        "comments": "true",  
        "object_list": [  
          {"owner": "SELECT_AI_USER", "name": "table1"},  
          {"owner": "SELECT_AI_USER", "name": "table2"},  
          {"owner": "SELECT_AI_USER", "name": "table3"}  
        ]}');  
end;
```

Select AI supports **conversations** for chatbot-like behavior

Enable conversational LLM interaction

Select AI remembers your last 10 prompts and adds them to current prompt

Keep the conversation going with 'chat'

- Start your chat
- Ask follow-up questions

Refine your data exploration with 'runsql', 'showsql', 'narrate', 'explainsql'

- Ask a question
- Review the response
- Follow up with more questions to refine, clarify, or get more info

Conversations using 'chat'

Enable conversational LLM interaction

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE (
    profile_name => 'OPENAI2',
    attributes =>
      '{"provider": "openai",
       "credential_name": "OPENAI_CRED",
       "model": "gpt-4",
       "conversation": "true"}');
END;
```

%sql

```
SELECT AI chat What is an autonomous database
```



RESPONSE

An autonomous database is a cloud-based technology designed to automate many of the routine tasks required to manage a database, such as tuning, patching, backups and updates. This type of database uses machine learning algorithms to automate database tuning, security, backups, updates, and other routine management tasks traditionally performed by database administrators. This allows for reduced costs, increased efficiency and accuracy, and frees up human resources for more complex tasks.

%sql

```
SELECT AI chat When did it first release
```



RESPONSE

Oracle Autonomous Database was first announced in October 2017 and was officially launched by Oracle Corporation in February 2018.

Conversations for data

Enable conversational LLM interaction

Keep refining your question until you get the answer you need

- Ask a question
- review the response
- follow up with more questions to refine, clarify, or get more info

Example

- What are our total streams?
- Break that out by genre
- Add customer segment
- Keep the top 5 customer segments and genres by total views

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE (
    profile_name => 'your_profile',
    attributes =>
      '{"provider": "azure",
       "azure_resource_name": "your-resource",
       "azure_deployment_name": "your-deployment",
       "credential_name": "AI_CREDENTIAL",
       "comments": "true",
       "conversation": "true",
       "object_list": [
         {"owner": "moviestream", "name": "GENRE"},
         {"owner": "moviestream", "name": "CUSTOMER"},
         {"owner": "moviestream", "name": "PIZZA_SHOP"},
         {"owner": "moviestream", "name": "STREAMS"},
         {"owner": "moviestream", "name": "MOVIES"},
         {"owner": "moviestream", "name": "ACTORS"}
       ]
      }';
);
END;
```

Enable chatbot-like conversations

```
sql-scripts > conversations.sql > ...
35  -- Have a Conversation with your data
36  -- Ask an initial question and then build on that - just like a normal conversation
37
38  -- Begin with a high level question
39  select ai Start with our total streams;
40
41  -- Give me more details
42  select ai Break that out by genre;
43
44  -- Which customer segments are watching the genres?
45  select ai Add customer segment;
46
```

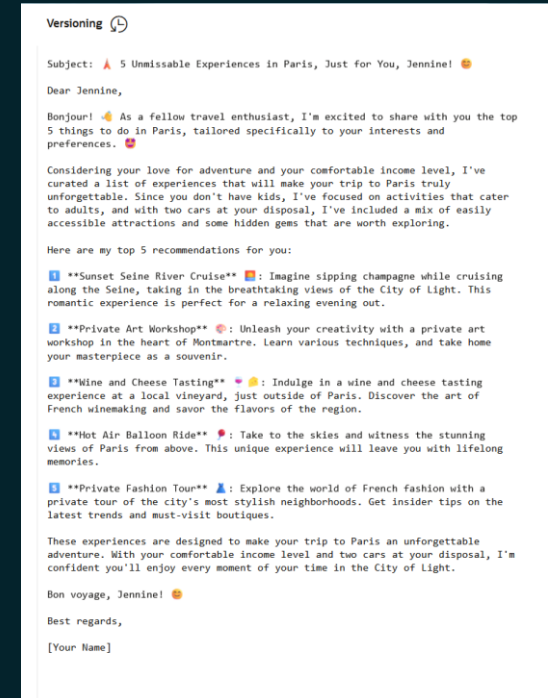
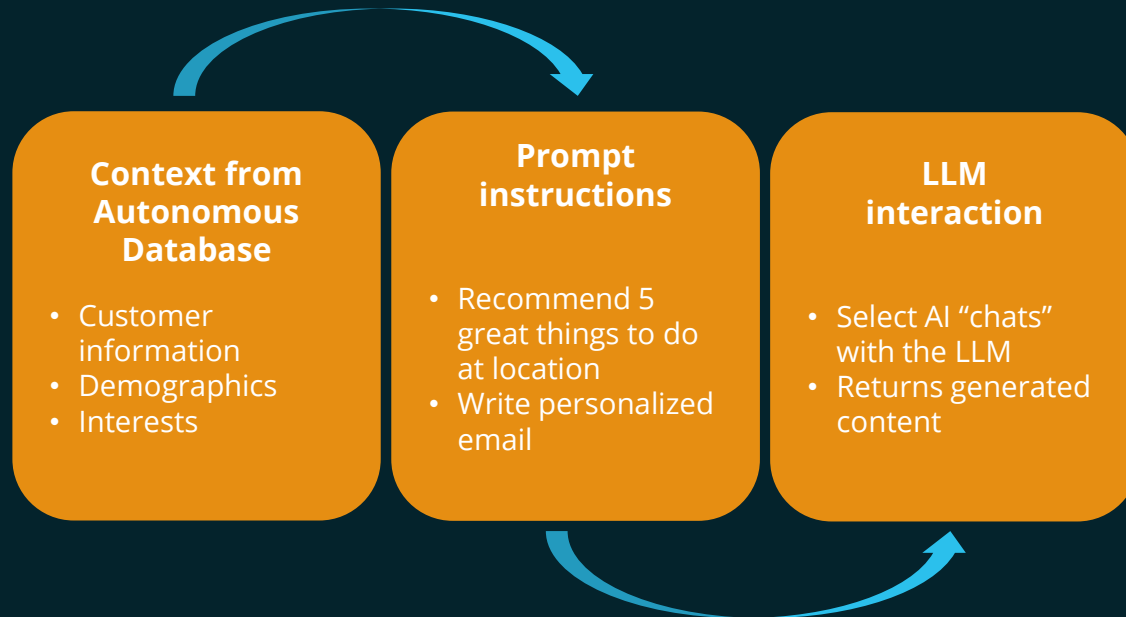
PROBLEMS 1 OUTPUT DEBUG CONSOLE TERMINAL PORTS QUERY RESULT SCRIPT OUTPUT

All rows fetched: 24 in 0.160 seconds

	GENRE_NAME	TOTAL_STREAMS
1	Documentary	4344
2	Family	126255
3	Mystery	34161
4	Adventure	255163

Enhance GenAI with context from your database

Generate an email to a customer with activity recommendations for a customer-specific destination



Compelling, personalized promotion

Create structured application-specific prompts for the LLM

Query task and customer data

→ combine tasks and contextual data

```
SELECT task,  
       task_rules,  
       last_name,  
       first_name,  
       location,  
       age,  
       gender,  
       has_kids,  
       num_cars,  
       income_level,  
       dog_owner  
FROM v_customer,  
     genai_project_task  
WHERE customer_id = 1  
AND    task_id = 3
```

Tabular result

TASK	TASK_RULES	LAST_NAME	FIRST_NAME	LOCATION	AGE	...	DOG_OWNER
Pick 5 great things to do at the location	<ol style="list-style-type: none">1. Encourage the customer to do these things. Really sell them as to why it would be right for them2. Consider all the information about the customer that's provided, including family and dog situation, whether they have a car and income3. Format the result with emojis and make it fun	Mouly	Jennine	Paris, Ile-De-France	37	...	NO

Create structured application-specific prompts for the LLM

Query task and customer data

```
SELECT JSON_OBJECT(*)
FROM (SELECT task,
  task_rules,
  last_name,
  first_name,
  location,
  age,
  gender,
  has_kids,
  num_cars,
  income_level,
  dog_owner
FROM v_customer,
  genai_project_task
WHERE customer_id = 1
AND task_id = 3)
```

JSON document

```
{
  "task": "Pick 5 great things to do at the location",
  "task_rules": "
1. Encourage the customer to do these things. Really sell them as to why
it would be right for them
2. Consider all the information about the customer that's provided,
including family and dog situation, whether they have a car and income
3. Format the result with emojis and make it fun",
  "last_name": "Mouly",
  "first_name": "Jennine",
  "location": "Paris, Île-De-France",
  "age": 37,
  "gender": "Female",
  "has_kids": "NO",
  "num_cars": 2,
  "income_level": "F: Above 110,000",
  "dog_owner": "NO"
}
```

Create structured application-specific prompts for the LLM

Use GENERATE to make 'chat' request

```
DECLARE
  v_response VARCHAR2(4000);
BEGIN
  v_response := DBMS_CLOUD_AI.GENERATE(prompt => 'Generate email using the JSON specification
{"task": "Pick 5 great things to do at the location",
"task_rules": "
  1. Encourage the customer to do these things. Really sell them as to why it would be right for them
  2. Consider all the information about the customer that's provided, including family and dog situation,
  whether they have a car and income
  3. Format the result with emojis and make it fun",
"last_name": "Mouly",
"first_name": "Jennine",
"location": "Paris, Île-De-France",
"age": 37,
"gender": "Female",
"has_kids": "NO",
"num_cars": 2,
"income_level": "F: Above 110,000",
"dog_owner": "NO"}',
profile_name => 'OCI_GENAI_PROFILE',
action      => 'chat');
  DBMS_OUTPUT.PUT_LINE(v_response);
END;
```

Response from LLM

Versioning ↻

Code Editor icons: `</>`, `🔍`, `📄`

Subject: 🇫🇷 Your Personalized Parisian Adventure Awaits, Jennine! 🎉🍷

Dear Jennine Mouly,

Bonjour! 🇫🇷 We're thrilled to hear about your upcoming trip to the City of Lights, Paris, Île-De-France. Given your adventurous spirit, we've curated a list of five fantastic experiences that we believe will make your trip unforgettable. ✨

1. 🗣️ The Louvre: As an art lover, you can't miss the world's largest art museum. The Louvre is home to thousands of works of art, including the Mona Lisa. With your income level, you can afford a private tour to avoid the crowds and enjoy the art in peace. 🏠
2. 🍷 Wine Tasting in Montmartre: Paris is synonymous with fine wine. A wine tasting tour in the charming district of Montmartre will allow you to savor the best French wines. And since you have two cars, why not buy a few bottles to take home? 🚗
3. 🛍️ Shopping on the Champs-Élysées: This famous avenue is a paradise for fashionistas. With your income, you can indulge in some luxury shopping. Don't forget to visit the iconic Louis Vuitton flagship store! 🛍️
4. 🍴 Gourmet Dining: Paris is a foodie's dream. With your budget, you can enjoy a meal at a Michelin-starred restaurant like Le Jules Verne, located in the Eiffel Tower. Imagine dining with a view of the city lights! 🌃
5. 🚤 Seine River Cruise: End your day with a romantic cruise on the Seine River. The city's landmarks look even more beautiful from the water, especially at sunset. 🌅

We hope these suggestions inspire you to make the most of your Parisian adventure. Remember, Paris is not just a city, it's a world. Enjoy every moment, Jennine! 🌍

Safe travels,
[Your Name]

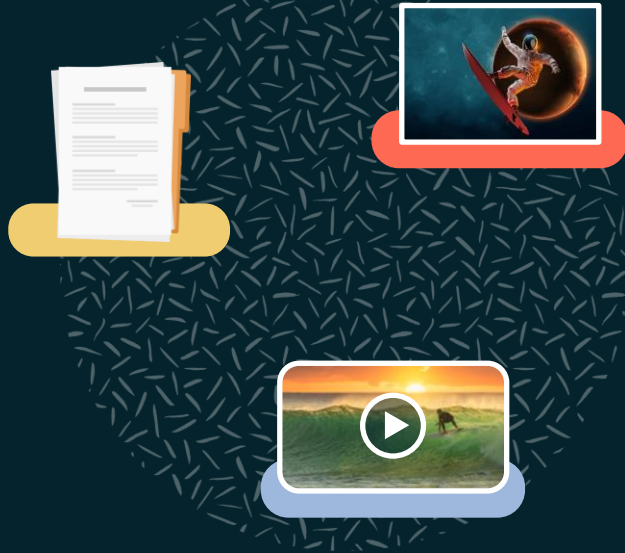
PL/SQL procedure successfully completed.

Retrieval Augmented Generation (RAG)

Simplifies RAG workflow so any database user can use it!

- ✓ Give the LLM new knowledge without fine-tuning
- ✓ Use natural language for semantic similarity search and LLM response generation
- ✓ Seamless integration with Oracle AI Vector Search
- ✓ Automate orchestration steps with fully managed Vector Index pipeline for new data

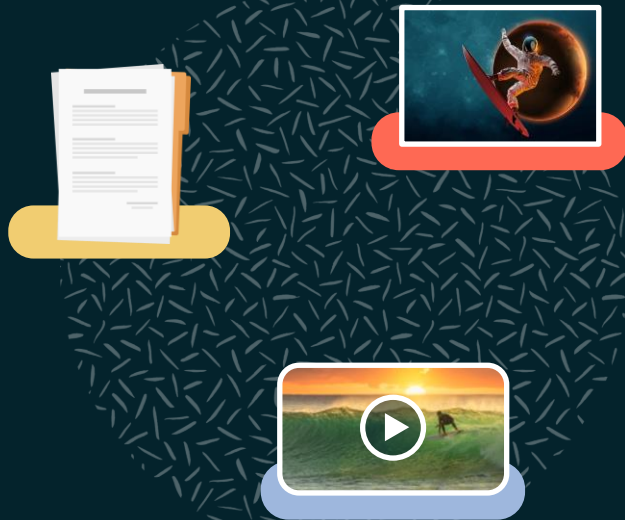




Databases are great at querying business data that is stored as strings, numbers, and dates

They have not been good at querying *unstructured data* such as text, images, audio, and videos

Semantic similarity search involves querying related content using vectors



Vector

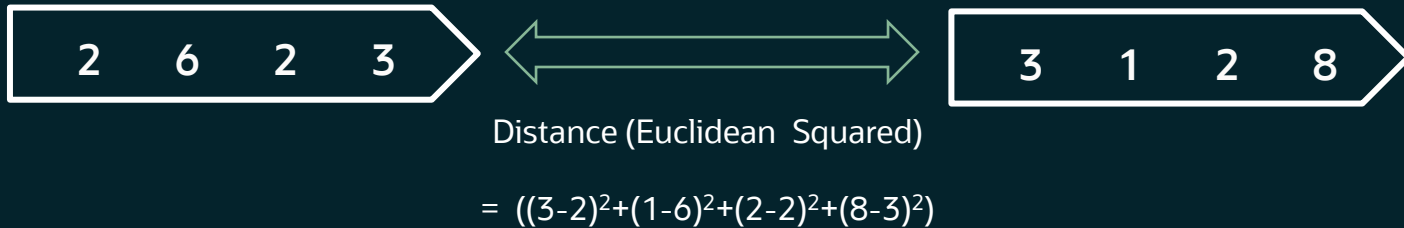
A *vector* is a sequence of numbers, called dimensions, that represent the semantic content of a document, image, audio, or video

Vectors represent the *semantic content* of data, not the underlying words or pixels

Deep learning *transformers* (or *embedding models*) generate vectors

The terms *vector* and *embedding* are often used interchangeably

The main operation on vectors is the mathematical distance between them



Oracle Database supports multiple mathematical distance functions

Retrieval Augmented Generation (RAG)

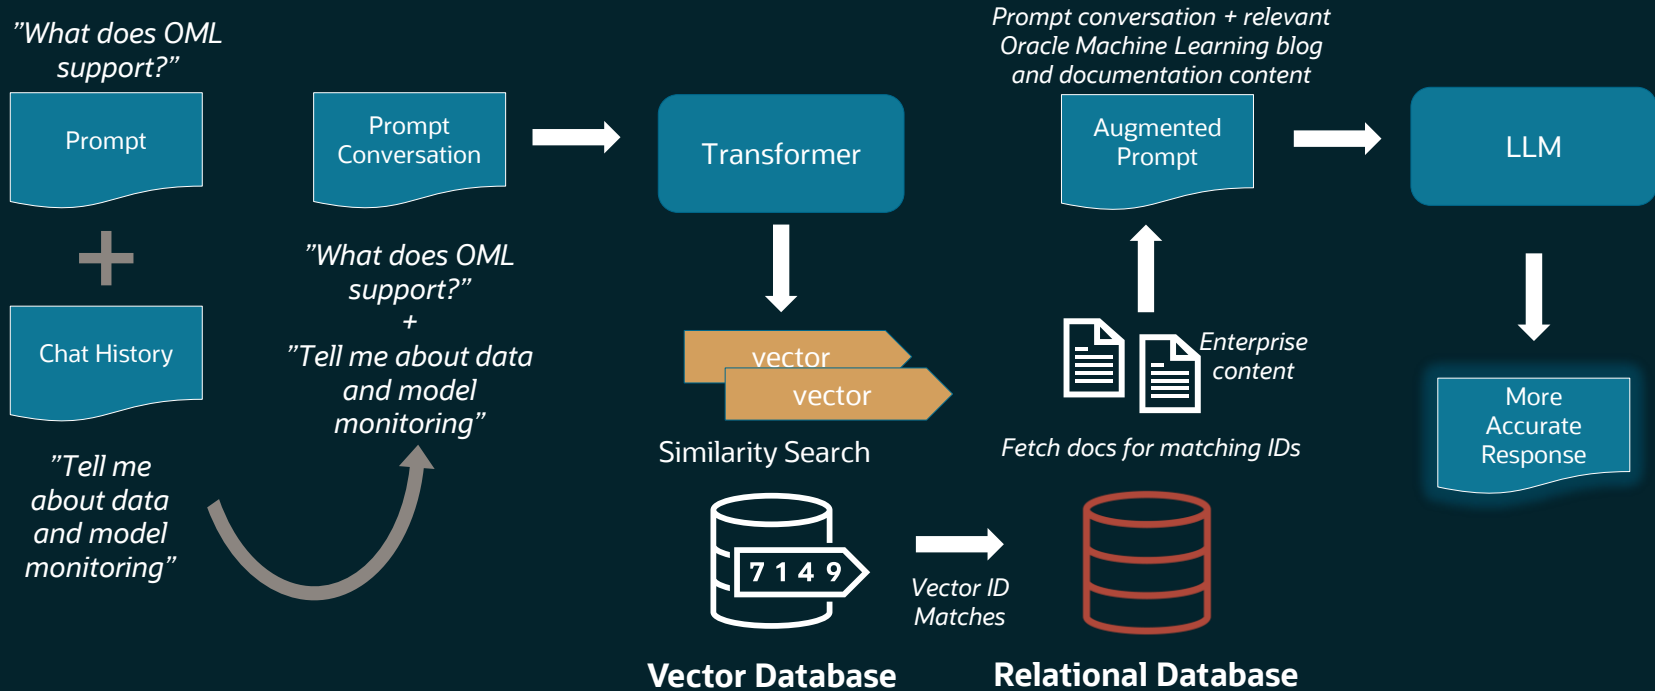


A technique that uses **private database** content to **augment user-provided prompts** using semantic similarity search with LLMs

RAG enables LLMs to use business data to produce better and more accurate responses and not fine-tune LLMs using that data, which may introduce security concerns

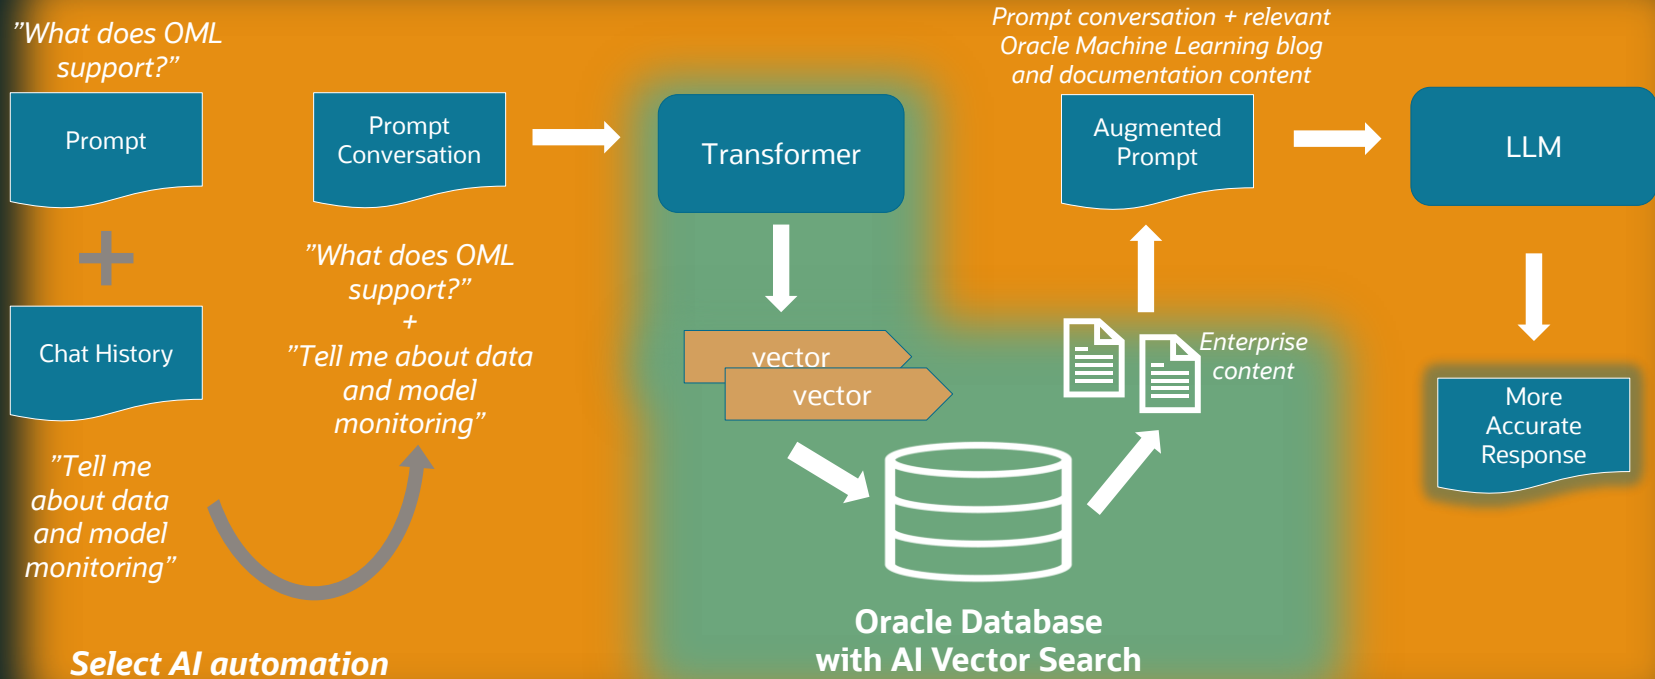
RAG pipeline example

Search Oracle Machine Learning blogs to answer question



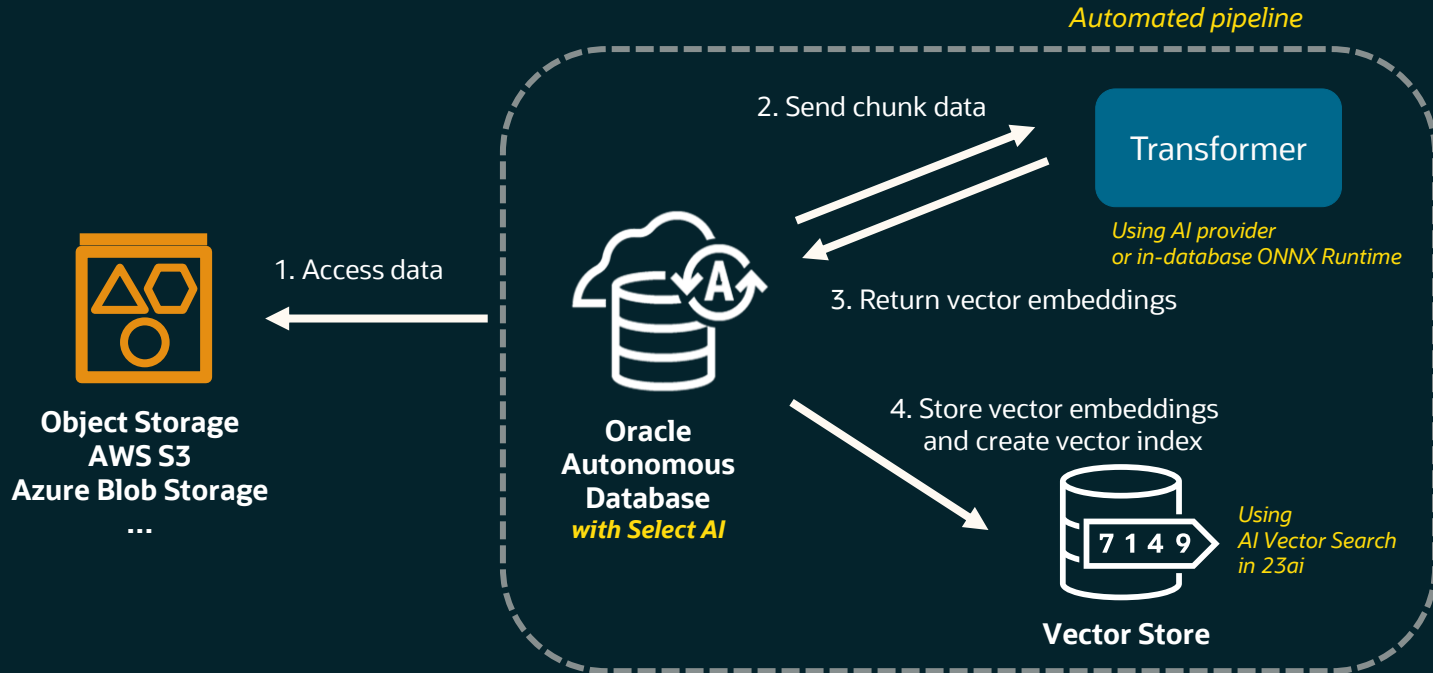
RAG pipeline example

Search Oracle Machine Learning blogs to answer question



Select AI Retrieval Augmented Generation (RAG)

Step 1: Create your vector content...automated by Select AI



Create a vector index

Example

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name => 'OPENAI_TEXT_TRANSFORMER',
    attributes   => '{"provider": "openai",
                    "credential_name": "OPENAI_CRED",
                    "embedding_model": "text-embedding-ada-002" }');
END;
```

Specify your AI provider and embedding model

```
BEGIN
  DBMS_CLOUD_AI.CREATE_VECTOR_INDEX(
    index_name => 'MY_VECTOR_INDEX',
    attributes => '{"vector_db_provider": "oracle",
                  "location": "https://my_namespace/my_bucket/my_data_folder",
                  "object_storage_credential_name": "OCI_CRED",
                  "profile_name": "OPENAI_TEXT_TRANSFORMER",
                  "vector_dimension": 1536,
                  "vector_distance_metric": "cosine",
                  "chunk_overlap": 128,
                  "chunk_size": 1024}');
END;
```

Create the vector index

Create a RAG-enabled AI Profile for use with 'narrate' and 'runsql'

Example

```
BEGIN
  DBMS_CLOUD_AI.CREATE_PROFILE(
    profile_name => 'OPENAI_GPT',
    attributes   => '{"provider": "openai",
                    "credential_name": "OPENAI_CRED",
                    "vector_index_name": "MY_VECTOR_INDEX",
                    "temperature": 0.2,
                    "max_tokens": 4096,
                    "model": "gpt-4o",
                    "embedding_model": "text-embedding-ada-002",
                    "enable_sources": true }');
END;
```

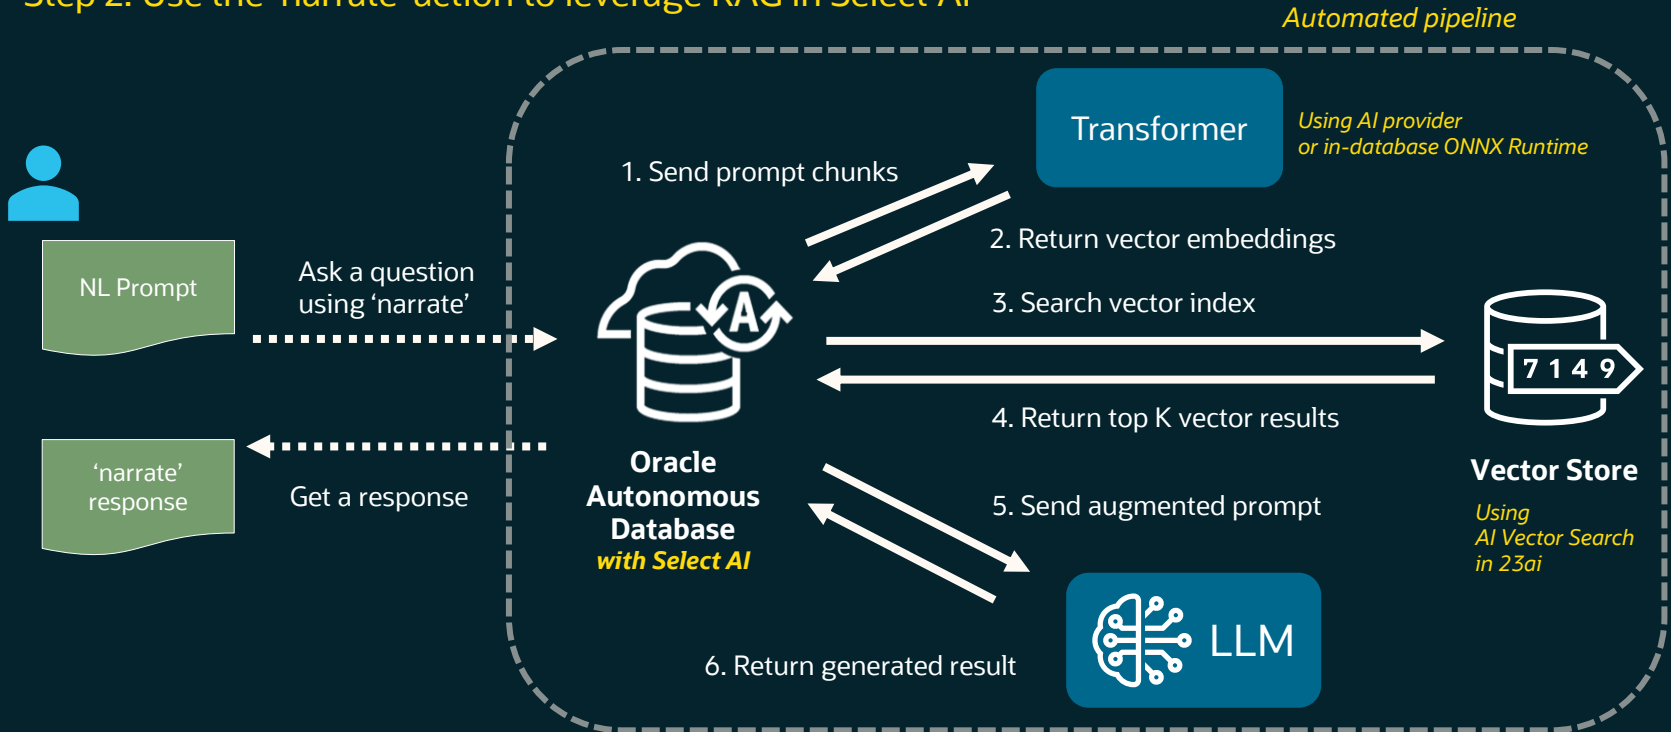
Specify vector index

Specify LLM

Specify the transformer

Select AI Retrieval Augmented Generation (RAG)

Step 2: Use the 'narrate' action to leverage RAG in Select AI



Example using 'narrate' on Oracle Machine Learning blogs

```
EXEC DBMS_CLOUD_AI.SET_PROFILE('OPENAI_GPT');
```

```
SELECT AI NARRATE What type of monitoring is enabled in Oracle Machine Learning
```

RESPONSE

Oracle Machine Learning enables two types of monitoring: Data Monitoring and Model Monitoring.

Data Monitoring is a no-code user interface that provides insight into how enterprise data evolves over time. It helps to identify data drift and gain insight into individual data features and their interactions. This tool is essential for maintaining data integrity for enterprise applications and dashboards.

Model Monitoring, on the other hand, helps maintain the accuracy of machine learning models and the effectiveness of applications. It detects concept drift and quality drift, tracking and reporting changes in model prediction patterns and accuracy. This helps to know when models need to be rebuilt or when other causes such as data drift need to be investigated.

Sources:

- [Announcing-OML-Data-Monitoring-User-Interface-on-Autonomous-Database.txt](https://objectstorage.../select_ai_rag_demo_1/o/Announcing-OML-Data-Monitoring-User-Interface-on-Autonomous-Database.txt)
(https://objectstorage.../select_ai_rag_demo_1/o/Announcing-OML-Data-Monitoring-User-Interface-on-Autonomous-Database.txt)
- [Announcing-OML-Monitoring-on-Autonomous-Database.txt](https://objectstorage.../select_ai_rag_demo_1/o/Announcing-OML-Monitoring-on-Autonomous-Database.txt) (https://objectstorage.../select_ai_rag_demo_1/o/Announcing-OML-Monitoring-on-Autonomous-Database.txt)
- [Simplify-your-model-monitoring-and-MLOps-with-OML-Model-Monitoring-UI.txt](https://objectstorage.../select_ai_rag_demo_1/o/Simplify-your-model-monitoring-and-MLOps-with-OML-Model-Monitoring-UI.txt)
(https://objectstorage.../select_ai_rag_demo_1/o/Simplify-your-model-monitoring-and-MLOps-with-OML-Model-Monitoring-UI.txt)

Example using 'runsql'

```
SELECT AI RUNSQL What type of monitoring is enabled in Oracle Machine Learning
```

DATA	SOURCE	URL	SCORE
----- We are pleased to announce the availability of Oracle Machine Learning (OML) Monitoring as part of OML Services on Oracle Autonomous Database. With OML Monitoring, you can be alerted to issues in both data and native in-database model quality. In data-driven enterprises, you need to know if there	----- Announcing-OML-Monitoring-on-Autonomous-Database.txt	----- <a href="https://<bucket>/Announcing-OML-Monitoring-on-Autonomous-Database.txt">https://<bucket>/Announcing-OML-Monitoring-on-Autonomous-Database.txt	----- 0.26

• • •

Select AI with RAG Demonstration

Synthetic Data Generation (SDG)

Realistic test data generated by LLMs

Support multiple use cases

- Jumpstart a new project

- Validate the user experience

- Integrate with Autonomous Database Metadata Clone for test-dev scenarios

Simple and Scalable



Example of synthetic data generation for **single table**

```
-- Create new or use existing table

-- Generate rows

BEGIN
  DBMS_CLOUD_AI.generate_synthetic_data(
    profile_name => 'GENAI',
    object_name  => 'Director',
    owner_name   => 'SELECT_AI_USER',
    record_count => 5,
    user_prompt  => 'born after 1950'
  );
END;
```

```
-- Query the table to see results

SQL> SELECT * FROM SELECT_AI_USER.Director

DIRECTOR_ID NAME
-----
1 John Smith
2 Emily Chen
3 Michael Brown
4 Sarah Taylor
5 David Lee

-- Or ask select ai to show the results
SQL> SELECT AI how many directors are there

NUMBER_OF_DIRECTORS
-----
5
```

Example of synthetic data generation for **multiple tables**

```
-- Create new/use existing tables with
-- optional referential integrity constraints

BEGIN
  DBMS_CLOUD_AI.generate_synthetic_data(
    profile_name => 'GENAI',
    object_list =>
      '[{"owner": "SELECT_AI_USER",
        "name": "Director", "record_count":5},

        {"owner": "SELECT_AI_USER",
        "name": "Movie_Actor", "record_count":5},

        {"owner": "SELECT_AI_USER",
        "name": "Actor", "record_count":10},

        {"owner": "SELECT_AI_USER",
        "name": "Movie", "record_count":5,
        "additional_prompt":
        "all movies are released in 2009"}]'
  );
END;
```

```
-- Query the table to see results
```

```
SQL> SELECT * FROM SELECT_AI_USER.Movie
```

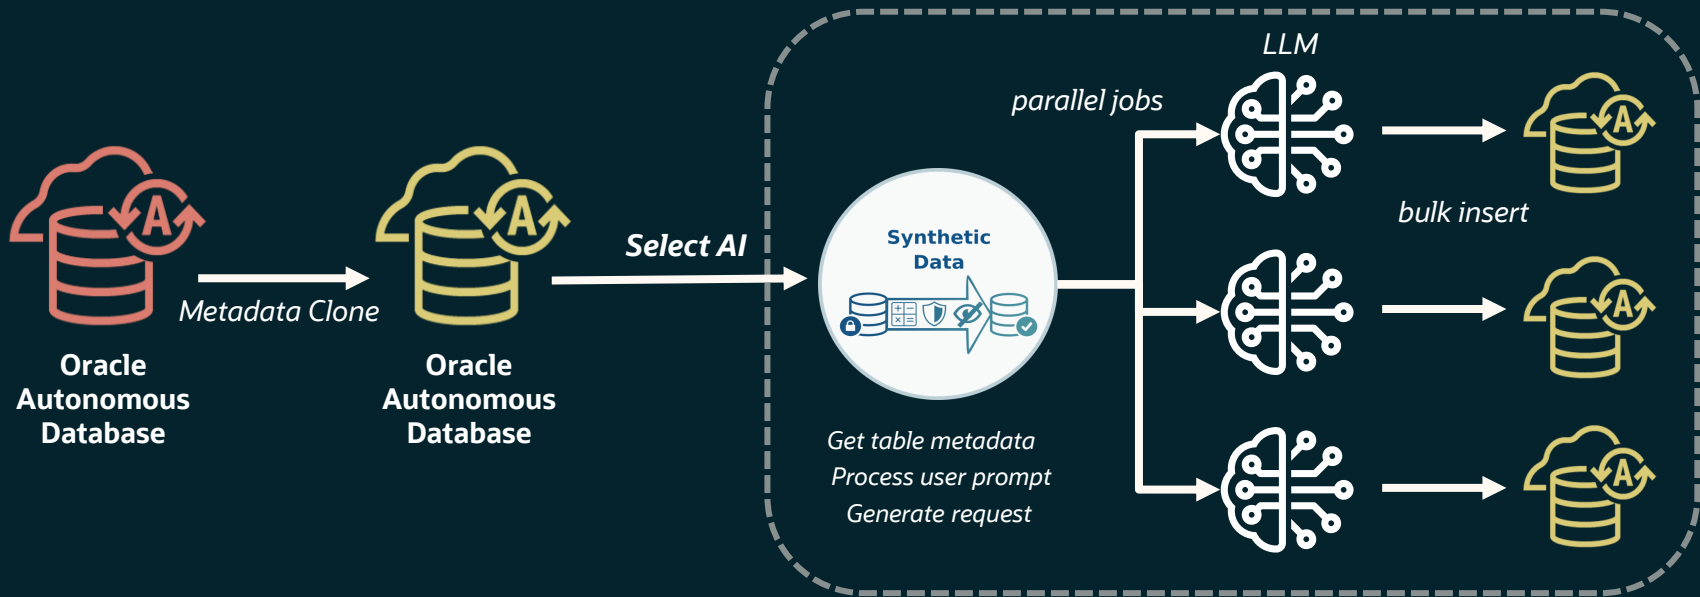
MOVIE_ID	TITLE	RELEASE_D	GENRE	DIR_ID
1	The Dark Knight	15-JUL-09	Action	8
2	Inglourious Basterds	21-AUG-09	War	3
3	Up in the Air	04-SEP-09	Drama	6
4	The Hangover	05-JUN-09	Comedy	1
5	District 9	14-AUG-09	Sci fi	10

```
-- Or ask select ai to show the results
```

```
SQL> SELECT AI how many actors are there
```

```
Number of Actors
-----
                10
```

Metadata Clone + Synthetic Data



Synthetic Data Generation Workflow

Select AI New Features

Recent New Features

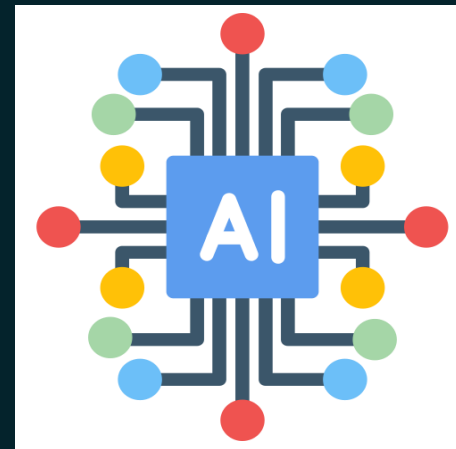
Support on-premises, non-ADB Oracle Database 23.7
(N12SQL only)

Use ONNX-format transformers from
in-database ONNX Runtime

```
"embedding_model": "database:ALL_MINILM_L12_V2"
```

AI Providers

- Google Gemini
- Anthropic Claude
- Hugging Face, e.g. Mistral, DeepSeek, etc.



Select AI Roadmap

Roadmap

Support on-premises, non-ADB Oracle Database 23.8

(with RAG and SDG)

Provide a conversation management API

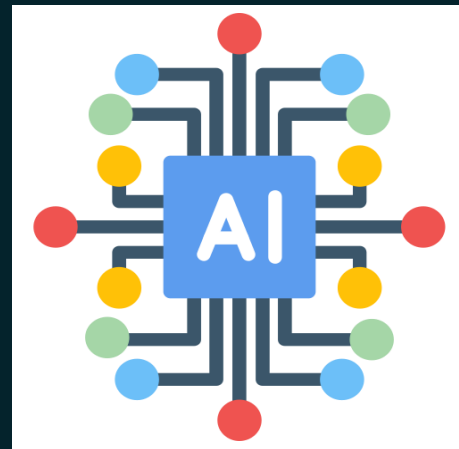
NL2SQL

RAG

AI Profile

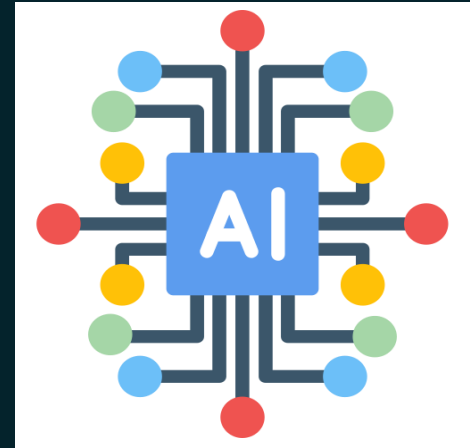
AI Providers

Actions



Roadmap: NL2SQL

- ✓ Support 23ai **ANNOTATIONS** on tables/columns in metadata
- ✓ Augment metadata with **foreign key constraints**
- ✓ Enable **case-insensitive** WHERE clause for better query responses
- ✓ **Historical view** of Select AI prompts and responses
- ✓ Support **Analytic Views**
- ✓ **Handle large schema** by automatically **detecting object list** for prompt metadata
- ✓ Collect feedback/corrections on generated SQL to **improve prompt tuning**



Roadmap: Handle large schema metadata

Automates/optimizes selection of schema object (tables/views) based on user prompt

Can significantly reduce metadata volume sent to LLM

Leverages AI Vector Search with transformer defined in AI profile

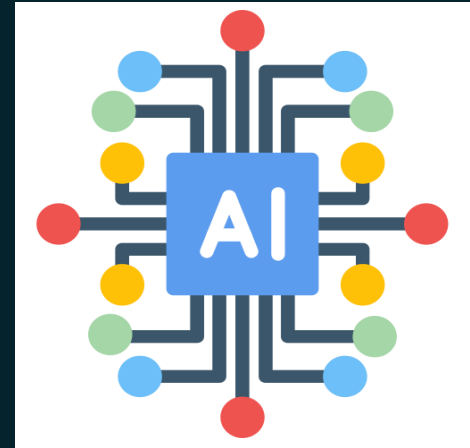
Process

- Vectorize schema metadata of all schema objects
- Periodically update vector data (per settable `refresh_rate`) to reflect schema changes
- Based on the vector similarity, retrieve metadata for top K (default `match_limit` 10) most relevant objects
- Adjust `similarity_threshold` to determine most contextually similar objects

Enabled using AI Profile attribute `"object_list_mode": "automated"`

Roadmap: RAG

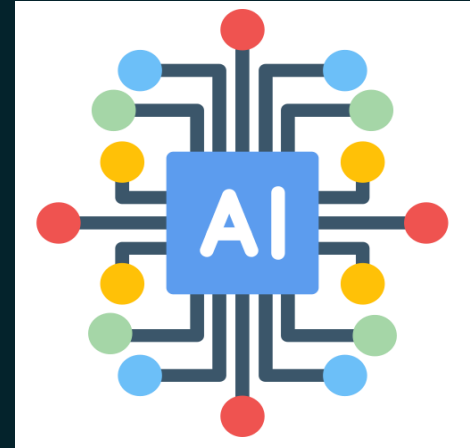
- ✓ Process PDF, XLS, and PPT files for vector index creation
- ✓ Include chunk offsets in RAG sources listing
- ✓ Multi-modal LLM and transformer supporting images and text
- ✓ Hybrid vector index support
- ✓ Graph RAG
- ✓ Enable vector index creation from database table columns
- ✓ Support multiple document sources in same AI profile
- ✓ Additional document sources from directory objects (DBFS or NFS), Google Drive, and MS SharePoint



Roadmap: AI Profile

Settings to:

- ✓ Specify conversation **chat history length** and **retention period**
- ✓ **Restrict tables used** to those explicitly listed
- ✓ Enable using a single AI Profile for both **object list for NL2SQL** and **vector index with RAG**



Roadmap: Conversation Management API

Supports building ChatGPT-like chatbots

Persist conversations across sessions

Provide manual or LLM-generated conversation titles

Set the conversation retention period
(# days before deletion)

Facilitates use within Oracle APEX

API

- **RESET_CONVERSATION** ()
Remove history from current conversation (start over)
- **BEGIN_CONVERSATION** (name)
Start new named conversation
- **DELETE_CONVERSATION** (name)
Remove named conversation
- **LOAD_CONVERSATION** (name)
Load named conversation and continue using
- **LIST_CONVERSATIONS** ()
Return list of existing named conversations
- **SELECT * FROM V\$** view
Return rowset of existing named conversations

Roadmap: AI Providers

- ✓ Support privately hosted LLMs using Ollama over [private endpoints](#)

In AI Profile

```
"provider": "ollama"
```

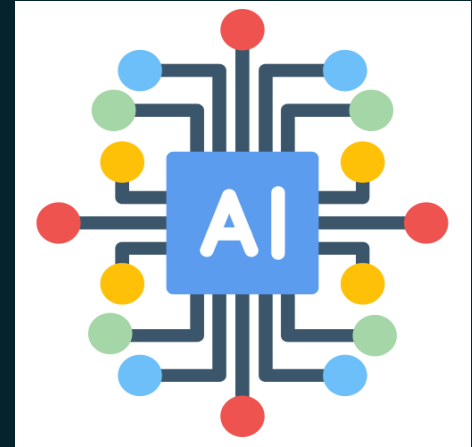
```
"provider_endpoint": "myhost.oraclevcn.com"
```

and then...

```
alter database property
```

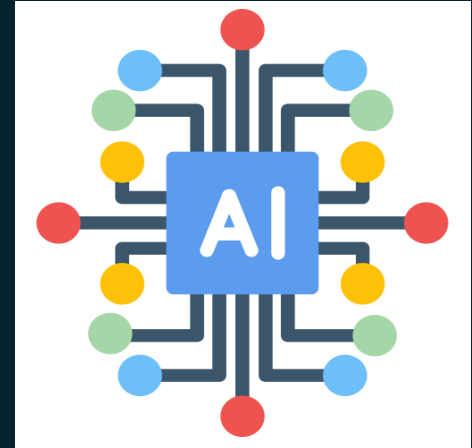
```
set route_outbound_connections = 'ENFORCE_PRIVATE_ENDPOINT'
```

- ✓ ADB-provided [LLM](#)
- ✓ AWS [Bedrock](#) and [OpenAI-compatible](#) providers



Roadmap: **Actions**

- ✓ 'summarize' – return text summarization of long-passage text
- ✓ 'feedback' – user feedback on NL2SQL results to improve future queries
- ✓ 'embedding' – return prompt embedding from transformer
- ✓ 'showprompt' – return augmented prompt sent to LLM
- ✓ 'admin' – support for database administration



Roadmap: Prompt tuning with 'feedback' action

-- SQL Text

```
SQL> select ai feedback for query  
"select ai showsql how many watch histories in  
total", use sum instead of count
```

```
SQL> select ai feedback for query  
"select ai showsql how many watch histories in  
total", the sql query generated is correct
```

-- SQL ID

```
SQL> select ai feedback  
use sum instead of count for sql_id 1v1z68ra6r9zf
```

```
SQL> select ai feedback  
sql query result is correct for sql_id 1v1z68ra6r9zf
```

-- Latest SELECT AI statement invoked (default)

```
SQL> select ai feedback  
use ascending sorting for ranking
```

```
SQL > select ai feedback  
the result is correct
```

```
DBMS_CLOUD_AI.FEEDBACK(profile_name IN VARCHAR2,  
                        sql_text     IN CLOB,  
                        feedback_type IN VARCHAR2,  
                        response      IN CLOB);
```

```
DBMS_CLOUD_AI.FEEDBACK(profile_name IN VARCHAR2,  
                        sql_id       IN VARCHAR2,  
                        feedback_type IN VARCHAR2,  
                        response      IN CLOB);
```

Roadmap: AI for ADMINs

Simplify database administration using AI

Troubleshoot SQL performance

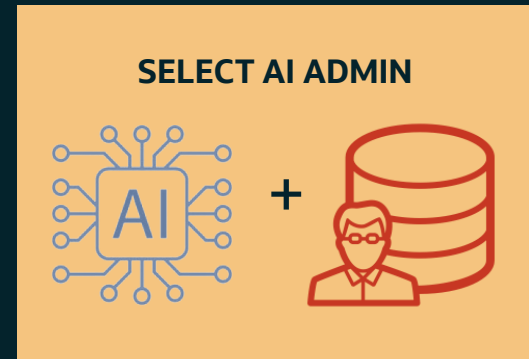
Identify Application bottlenecks

Monitor Database resources

Enable via new action 'admin'

LLM categorizes user question

Category determines specific prompt instructions



Roadmap: AI for ADMINs

```
SQL> SELECT AI ADMIN find any issues with queries on employee table
```

```
RESPONSE
```

```
-----  
The SQL statement `select * from employee` has an issue related to the Cost-Based Optimizer (CBO) system statistics. The observation indicates that the workload CBO statistics are not gathered, and the CBO is using default values. This can lead to suboptimal execution plans.
```

```
To address this, you should gather system statistics using the `DBMS_STATS.GATHER_SYSTEM_STATS` procedure. Here is an example of how you can do this:
```

```
BEGIN  
  DBMS_STATS.GATHER_SYSTEM_STATS('INTERVAL', interval => 60);  
END;  
/
```

```
This command gathers system statistics over a 60-minute interval. Adjust the interval as needed based on your workload and system activity.
```

For more information on Select AI...

Blogs

[Announcing Select AI with Retrieval Augmented Generation \(RAG\) on Autonomous Database](#)

[Announcing Select AI for Synthetic Data Generation](#)

[Introducing Select AI Natural Language to SQL on Autonomous Database](#)

[Announcing Select AI with Azure OpenAI Service on Autonomous Database](#)

[Announcing Select AI on Autonomous Database with OCI Generative AI Service](#)

[How to help AI models generate better natural language queries](#)

[Conversations are the next generation in natural language queries](#)

[Accelerate innovation with enterprise data, OCI Generative AI, and enhanced security](#)

Documentation

<https://docs.oracle.com/en-us/iaas/autonomous-database-serverless/doc/sql-generation-ai-autonomous.html>

LiveLabs

[Chat with Your Data in Autonomous Database Using Generative AI](#)

[Develop apps using GenAI, Autonomous Database, and React](#)

Videos

[Simplify Developing RAG Applications](#)

[Autonomous Database Speaks "Human" using Select AI](#)

Register for the Select AI Office Hours series



Getting Started with Select AI

