

Bringing AI to Data: Lessons and Experiences with Oracle AI Database

May 31, 2026


Tirthankar Lahiri

SVP, Mission Critical Data and AI Engines




AI Evolution

Agentic AI evolves AI from answers to actions



AI Evolution

Beyond Agentic AI?




*The time has come, the Database said,
To talk of many things
Of vectors, graphs, and icebergs deep,
And agents pulling strings.*


.. To be Continued ...

Whose Prompt is it Anyway?

How do you build a class-leading AI native enterprise database?




How do you build a class-leading AI native enterprise database?



Converged Architecture

Build best-of-breed data types and workloads,
engineered to work together

Using a unified engine with scale-up, scale-out
security and disaster recovery




Converged data architecture supports all key Workloads

Optimized algorithms
for each workload

Scale to meet any need

Even better for apps
using mixed workloads

NEW: Support for graph queries
and Iceberg Data Lake



All Workloads are Engineered to Work Together


Converged data architecture supports all key Data Types

Optimized for each type

Ideal for apps with a mix
of data types


For a new data type or workload,
add a SQL statement, not a DB

NEW: AI Vector type, Ultra-fast JSON,
ability to run MongoDB Apps



All Data Types are Engineered to Work Together


Converged architecture runs on a common mission-critical Data Engine



Synergy from common atomicity, parallelism, high availability, disaster recovery, and security

Across all workloads and types

NEW: Lock-free updates and transactional microservices



All Data Operations are Engineered to Work Together

Convergence Enables Unification

JSON Relational Duality and Property Graph Views **unify** key data models

Apps can access **the same data** as different types for different workloads


- No data duplication
- Not tied to one type and its limitations
- Perform analytics across all data shapes

Relational, JSON, and Graph unified on the same stored format

Access data as
JSON document

Access data as
Relational table

Access data
as Graph



Oracle Exadata Engineered System

The foundational building block for AI for data



Data-optimized hardware

Data intelligent software

Engineered for mission-critical



Latest CPUs and
RDMA networking

Exascale elastic
storage

SQL in storage with
vector search

Oracle Globally Distributed Database

Scale without limits, survive system failures and address data sovereignty regulations

Deployed by many top banks, telecoms, and manufacturers around the world



Single logical database to multiple physical databases



Unlimited Scalability



3-second failover




Data Sovereignty



Available on Exascale infrastructure
Multi-region RAFT replication
AI-ready with Vector Search and Agentic AI

How do you build a class-leading AI native enterprise database?



Whose Prompt is it Anyway?

*The time has come, the Database said,
To talk of many things
Of vectors, graphs, and icebergs deep,
And agents pulling strings.*

*Of prompts and skills and react chains,
Of grounding AI acts
Of why all shiny LLMs
Still need transactional facts.*

.. To be Continued ...



AI Vector Search

A **new** breakthrough technology to search documents, images, and other structured and unstructured data based on their **semantics**, rather than their **values**





AI Vector Search works by encoding the **semantics** of a document, image, video, or even relational data as a sequence of numbers, called a **vector**

Developers create a vector for a given piece of data by invoking their chosen deep-learning model (known as an **embedding model**)

Oracle Vector Database natively **stores** vectors and **compares** vectors to find objects with **similar semantic content**





AI Vector Search works by encoding the **semantics** of a document, image, video, or even relational data as a sequence of numbers, called a **vector**


Developers create a vector for a given piece of data by invoking their chosen deep-learning model (known as an **embedding model**)

Oracle Vector Database natively **stores** vectors and **compares** vectors to find objects with **similar semantic content**



An Example Business Scenario: **Support Incident Matching**


Support Incident



Note: Features are chosen by ML algorithms and are not as simple as shown here

An Example Business Scenario: **Support Incident Matching**

The vector for a support incident could be ...




Each dimension (number), represents a different feature of the support incident


Note: Features are chosen by ML algorithms and are not as simple as shown here



Support Incidents that are more **similar** have vectors **closer** to each other



Support Incidents that are more **similar** have vectors **closer** to each other



Searches on a **combination** of **business** and **semantic** data is more effective if both types of data are stored together



Business Database

One solution is to send your business data to a vector database continuously




Vector Database

Searches on a **combination** of **business** and **semantic** data is more effective if both types of data are stored together



Business Database

Since you don't know the question, you need to share additional data



Vector Database

The business data that is relevant to a question **varies widely**
Plus, dedicated vector databases are not as good at **searching** or **securing** business data

Better together

Business data and AI vectors

Uniquely combines sophisticated business data search with vector similarity search using simple SQL

No need to move and synchronize data, manage multiple products, etc.

Every mission-critical feature of Oracle Database works transparently with AI Vector Search

Oracle Database's robust security ensures compliance with corporate security standards

Allows AI Vectors to be used immediately in enterprise apps of any scale or criticality



JSON




7149



Oracle's Converged Data Architecture

SQL easily supports **AI vector search on unstructured data**

Find the top 10 matching incidents



Support Rep Jane Doe
jane@doe.com

Laptop Gen 32


- Severity 1
- Spontaneous reboot
- Resolved
- Applied OS Update 42



```
SELECT ...  
FROM   Support_Incidents  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY;
```

SQL easily supports **AI vector search on unstructured data**


Find the top 10 matching incidents
for laptops



```
SELECT ...  
FROM   Support_Incidents S, Products P  
WHERE  P.type = 'Laptop'  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY;
```

SQL easily supports **AI vector search on unstructured data**

Find the top 10 matching incidents
for laptops
by customers in India



```
SELECT ...  
FROM   Support_Incidents S, Products P, Customers.C  
WHERE  S.prod_id = P.id AND S.cust_id = C.id  
AND    P.type = 'Laptop' AND C.country = 'India'  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY;
```


Enables combining AI vector search on **unstructured** Support Data with **value-based search** on **structured data** for Products and Customers

Combines incident, customer and product data, and AI search in a few lines of SQL

A single integrated solution, all data fully consistent

Any developer can learn to use it in 5-minutes

Find the top 10 matching incidents for laptops
by customers in India



```
SELECT ...  
FROM   Support_Incidents S, Products P, Customers.C  
WHERE  S.prod_id = P.id AND S.cust_id = C.id  
AND.   P.type = 'Laptop' AND C.country = 'India'  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY;
```

Vector Embedding Generation | **Four Alternatives**

1

Use
Pre-created
embeddings

2

Use an external
embedding
cloud-service

3

Use an external
embedding
library in the mid-
tier

4

Use a database
resident
embedding model

Vector Embedding Generation **inside the database**

4

Use a database resident embedding model

Generate embeddings using the **VECTOR_EMBEDDING()** SQL function using an imported **ONNX** embedding model, so that no data leaves the database

```
-- import onnx embedding model
DBMS_VECTOR.load_onnx_model(
  directory => <database directory>
  file_name => 'my_embed_model.onnx'
  model_name => 'embed-model',
  metadata => <source>
);

-- generate vectors from support incidents
SELECT
  VECTOR_EMBEDDING(embed-model USING incident_data)
FROM Support_Incidents;
```


An exhaustive search
for top-K matches
will be 100% accurate
but slow as data
volumes grow

New vector indexes trade-off some search
accuracy for 100x speed up

Vector Indexes | Neighbor Graph Vector Index

Graph-based index where vertices represent vectors and edges between vertices represent *similarity*

In-Memory only index - highly efficient for both accuracy and speed



Graph Vector Index (e.g.
HNSW Index)



```
CREATE VECTOR INDEX support_idx  
ON Support_Incidents(incident_vector)  
ORGANIZATION INMEMORY NEIGHBOR GRAPH  
DISTANCE COSINE  
TARGET ACCURACY 90;
```

Specify default accuracy (recall) for index


Vector Indexes | Neighbor Partition Vector Index

Partition-based index with vectors clustered into table partitions based on *similarity*

Efficient scale-out index for unlimited data size



```
CREATE VECTOR INDEX support_idx
ON Support_Incidents(incident_vector)
ORGANIZATION NEIGHBOR PARTITIONS [LOCAL]
DISTANCE COSINE
TARGET ACCURACY 90;
```



Partition Vector Index (e.g.
IVF_FLAT index)

Specify default accuracy (recall) for index

Vector Indexes | Included Columns

Included Columns are **additional columns stored inside the vector index**

- Accelerates hybrid queries
- Using the included columns avoids the join with the base table

Query with Filter

```
  
SELECT id, house_photo  
FROM houses_for_sale  
WHERE city = 'San Francisco' --> use included column  
ORDER BY VECTOR_DISTANCE(house_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY
```

Houses for Sale Table


```
  
CREATE TABLE houses_for_sale  
(id            number,  
 price         number,  
 city          varchar2(400),  
 house_photo   blob,  
 house_vector  VECTOR(768, FLOAT32))
```

Index w/ included column


```
  
CREATE VECTOR INDEX HOUSE_IDX  
ON houses_for_sale(house_vector)  
INCLUDE (city) --> included column  
ORGANIZATION NEIGHBOR PARTITIONS;
```

Vector Indexes are always **Transactionally Consistent**

Neighbor Partition
Vector Index



Neighbor Graph
Vector Index



Oracle's AI Vector Search Indexes
maintain transactional consistency with
DML activity

INCIDENT_ID	INCIDENT_TEXT	INCIDENT_VECTOR
...	...	...
...	...	...
...	...	...



A new vector is inserted



Vector Search with Vector Indexes | Search with Filters



```
SELECT ...  
FROM   Support_Incidents S  
WHERE  S.incident_date >= SYSDATE - 7  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY  
TARGET ACCURACY 95
```



Support Rep: Jane Doe
jane@doe.com
Laptop Gen S2
• Severity 1
• Spontaneous reboot
• Resolved
• Applied OS Update 42

*Find the top 10 matching incidents
filed **within the past week***

Can override default index accuracy


Vector Search with Vector Indexes | Search with Filters

```
SELECT ...  
FROM   Support_Incidents S  
WHERE  S.incident_date >= SYSDATE - 7  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY  
TARGET ACCURACY 95
```

Find the top 10 matching incidents
filed *within the past week*

Can override default index accuracy

PRE-FILTER



- Filter first, then rank Top K
- Best for **high selectivity** filters


Vector Search with Vector Indexes | Search with Filters

```
SELECT ...  
FROM   Support_Incidents S  
WHERE  S.incident_date >= SYSDATE - 7  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY  
TARGET ACCURACY 95
```

Find the top 10 matching incidents
filed *within the past week*


Can override default index accuracy

PRE-FILTER



- Filter first, then rank Top K
- Best for **high selectivity** filters

IN-FILTER



- Filter during Top K search
- Best for **medium selectivity** filters


Vector Search with Vector Indexes | Search with Filters

```
SELECT ...  
FROM   Support_Incidents S  
WHERE  S.incident_date >= SYSDATE - 7  
ORDER BY VECTOR_DISTANCE(incident_vector, :search_vector)  
FETCH FIRST 10 ROWS ONLY  
TARGET ACCURACY 95
```

Find the top 10 matching incidents
filed *within the past week*


Can override default index accuracy

PRE-FILTER



- Filter first, then rank Top K
- Best for **high selectivity** filters

IN-FILTER



- Filter during Top K search
- Best for **medium selectivity** filters

POST-FILTER



- Rank Top K' first, then filter, keep Top K
- Best for **low selectivity** filters

Optimizer chooses the best search strategy



Every mission-critical feature
of Oracle Database works
transparently with AI Vectors


Allowing AI Vectors to be used
immediately in enterprise apps
of any scale or criticality



Real-Application Cluster



Parallel SQL



Transactions



Security



Analytics




Disaster Recovery


AI Vector Search | Scale-Out with Real Application Clusters

AI Vector search **transparently scales vector** processing across the compute nodes in a RAC cluster

With full data consistency



AI Vector Search | Scale-Out with offload to Exadata Smart Storage



Oracle's AI vector search can be **transparently offloaded** to intelligent Exadata storage for up to 30X faster AI Vector queries

Vector search queries are **automatically parallelized** across the storage servers

Each storage server independently computes the top-K matches enabling **faster Top-K processing**

Supports extreme scale environments with **thousands of concurrent AI vector searches**

Vector Indexes | Scale-Out w/ Globally Distributed Database


Vector Indexes can be scaled out using a globally distributed database

- Each vector index is **Shard-Local**, i.e., all data within the shard is indexed independently

Top K queries with a Shard Key **equality filter** are sent only to the **qualifying shard**

Top K queries without a Shard Key filter are **broadcast** to all shards

- Each shard performs its own Top K query, and the shard coordinator merges the results



Vectors On Ice


Unified AI Search across
Databases and Data Lakes

Now playing at an AI lakehouse near you



Vector Indexes on Iceberg Tables


Native support for vector data in Apache Iceberg Tables



Oracle can now **access vectors** stored in Apache Iceberg tables in addition to the AI Database




Oracle can now **create Vector Indexes** on vectors in Iceberg in addition to the AI Database




Oracle excels at **combined search** of Iceberg vectors and business data to answer business questions

AI Vector Search is used across the globe in a variety of industries



How do you build a class-leading AI native enterprise database?



Hybrid Vector Index | Overview


There are two approaches for searching document data:

- **Keyword Search:** Deterministic but does not fully capture document semantics
- **AI Vector Search:** Powerful semantic search technique, but not deterministic

Many real-world document searches require both techniques to be combined:

- E.g, Find support incidents about **Desktop Systems** that **contain "Reboot"** in the description
- Requires semantic match ("About Desktop Systems") plus keyword match ("containing Reboot")

A Hybrid Vector Index combines keyword search with vector search for improved accuracy



Hybrid Vector Index



```
CREATE HYBRID VECTOR INDEX  
support_incident_idx ON  
support_incidents(incident_desc)  
PARAMETERS('MODEL embed-model');
```

ONNX model that was imported
into the database earlier

Hybrid Vector Index | Overview


A Hybrid Vector Index combines keyword search with vector search for improved accuracy

Unlike a traditional index, a Hybrid Vector Index must be explicitly queried via specialized access APIs:

- APIs for search and for combining the results of searching by text and vectors



```
SELECT DBMS_HYBRID_VECTOR.SEARCH(  
  JSON('{"hybrid_index_name": "SUPPORT_HYBRID_IDX",  
    "search_fusion"      : "UNION",  
    "vector"             : {"search_text": "Desktop Systems"} ,  
    "text"               : {"contains"   : "reboot"},  
    "return"            : {"values"     : [ "score", "chunk_text" ]}  
  }')) ;
```




A support ticket card with a light beige background. At the top left are two speech bubble icons. At the top right is a circular profile picture of a woman. Below the profile picture, the text reads "Support Rep Jane Doe" and "jane@doe.com". The title "Laptop Gen 32" is in bold. Below the title is a list of items: "Severity 1", "Spontaneous reboot", "Resolved", and "Applied OS Update 42".

Support Rep Jane Doe
jane@doe.com

Laptop Gen 32

- Severity 1
- Spontaneous reboot
- Resolved
- Applied OS Update 42

How do you build a class-leading AI native enterprise database?



Whose Prompt is it Anyway?

*The time has come, the Database said,
To talk of many things
Of vectors, graphs, and icebergs deep,
And agents pulling strings.*

*Of prompts and skills and react chains,
Of grounding AI acts
Of why all shiny LLMs
Still need transactional facts.*


*We all praise models night and day,
Their fluency? It flies!
But who has trust? Who serves the truth?
The DB! It never lies!*

.. To be Continued ...

Data privacy is frequently managed at the application layer

With siloed, app-level controls


- You have to trust every one of your applications to do the right thing
- Privileged accounts introduce high risk
- Limited visibility into who can do/did what



AI greatly increases the risk of sensitive data exposure

New AI risks

- Agents can generate arbitrary SQL
- AI can bypass app-level controls, access limitations, and firewalls




Secure data at the source with Oracle AI Database


Oracle Deep Data Security

- Secure identity propagation
- Database-enforced row/column/cell access
- Policy-as-SQL

```
CREATE DATA GRANT employees_own_record  
AS  
SELECT ON hr.employees  
WHERE email = ORA_END_USER_CONTEXT.username  
TO employee_role;
```



How do you build a class-leading AI native enterprise database?



Whose Prompt is it Anyway?

*The time has come, the Database said,
To talk of many things
Of vectors, graphs, and icebergs deep,
And agents pulling strings.*


*Of prompts and skills and react chains,
Of grounding AI acts
Of why all shiny LLMs
Still need transactional facts.*

*We all praise models night and day,
Their fluency? It flies!
But who has trust? Who serves the truth?
The DB! It never lies!*

*For years, I ran your world unseen
Now, AI steals my glory?
It's when your bank account is cleaned,
It's then, that you'll be sorry!*

.. To be Continued ...

Retrieval Augmented Generation with Vectors enables users to talk to their data



Improve Enterprise Generative AI by **augmenting LLM prompts** with **private database content**

Enables LLMs to use business data to produce better and more relevant answers and keep business data secure


Retrieval Augmented Generation (RAG)

Oracle AI Database can pass the results of AI Vector search on private data to an LLM to directly answer the user's question



AI Vector Search powers Complete RAG AI pipeline

AI Vector Search and RAG in Oracle AI Database



Single SQL Statement for all of RAG

1. Vectorize the Question
using a SQL function



```
WITH TOP50 AS (SELECT plan_text FROM BENEFIT_PLANS
                ORDER BY VECTOR_DISTANCE(plan_vector,
                                           SELECT VECTOR_EMBEDDING(embedding_model
                                           USING :question_text AS data))

                FETCH FIRST 50 ROWS ONLY),

    LLM_PROMPT AS (SELECT ('Answer this question using the following context,\
                          QUESTION: ' || :question_text || ',CONTEXT: '
                          || LISTAGG(plan_text, CHR(10))) AS prompt_text
                   FROM TOP50)

SELECT DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT(prompt_text, json(:LLM_params))
       AS Answer
FROM   LLM_PROMPT;
```

Single SQL Statement for all of RAG

1. Vectorize the Question
using a SQL function

2. Find top
matching docs
using a SQL
subquery

```
WITH TOP50 AS (SELECT plan_text FROM BENEFIT_PLANS
                ORDER BY VECTOR_DISTANCE(plan_vector,
                                           SELECT VECTOR_EMBEDDING(embedding_model
                                                                    USING :question_text AS data))
                FETCH FIRST 50 ROWS ONLY),

LLM_PROMPT AS (SELECT ('Answer this question using the following context,\n
                        QUESTION: ' || :question_text || ',CONTEXT: '
                        || LISTAGG(plan_text, CHR(10))) AS prompt_text
                FROM TOP50)

SELECT DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT(prompt_text, json(:LLM_params))
       AS Answer
FROM   LLM_PROMPT;
```


Single SQL Statement for all of RAG

1. Vectorize the Question
using a SQL function

2. Find top
matching docs
using a SQL
subquery

3. Compose
prompt using a
SQL subquery

```
WITH TOP50 AS (SELECT plan_text FROM BENEFIT_PLANS
                ORDER BY VECTOR_DISTANCE(plan_vector,
                                           SELECT VECTOR_EMBEDDING(embedding_model
                                           USING :question_text AS data))
                FETCH FIRST 50 ROWS ONLY),

LLM_PROMPT AS (SELECT ('Answer this question using the following context,\n
                        QUESTION: ' || :question_text || ',CONTEXT: '
                        || LISTAGG(plan_text, CHR(10))) AS prompt_text
                FROM TOP50)

SELECT DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT(prompt_text, json(:LLM_params))
       AS Answer
FROM   LLM_PROMPT;
```

Single SQL Statement for all of RAG

1. Vectorize the Question
using a SQL function

2. Find top
matching docs
using a SQL
subquery

3. Compose
prompt using a
SQL subquery


4. Call the LLM
using a SQL
statement

```
WITH TOP50 AS (SELECT plan_text FROM BENEFIT_PLANS
                ORDER BY VECTOR_DISTANCE(plan_vector,
                                           SELECT VECTOR_EMBEDDING(embedding_model
                                           USING :question_text AS data))
                FETCH FIRST 50 ROWS ONLY),

LLM_PROMPT AS (SELECT ('Answer this question using the following context,\n
                        QUESTION: ' || :question_text || ',CONTEXT: '
                        || LISTAGG(plan_text, CHR(10))) AS prompt_text
                FROM TOP50)

SELECT DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT(prompt_text, json(:LLM_params))
       AS Answer
FROM   LLM_PROMPT;
```

How do you build a class-leading AI native enterprise database?



Agentic AI architected into Oracle AI Database



Agent Foundations

Core **data, memory, and execution** primitives for enterprise agents



Agent Builders

Tools and frameworks to build, connect, and operationalize agents



Agent Solutions

Domain-specific **patterns and apps** that accelerate time to value



Oracle AI Database Delivers Agent Memory

Built-in support for every layer of agent memory
within a single unified platform

<https://pypi.org/project/oracleagentmemory/>

Agent core memory

Foundational for AI agents

Agent memory is information and long-term context for agents:


- Derived from past interactions, outcomes, and data
- Memory helps agents improve over time

For example:

Tabular Data:

Name	Dept	Hire Date
Barb	Legal	25-May-88


Agent Fact: *"Barb is the longest serving employee in legal"*
 (changed on 16-May-26)



Oracle AI Database for Unified Agent Memory

Fragmenting memory across
many databases compromises:

- Efficiency, consistency, security, reliability




Oracle AI Database for Unified Agent Memory

Fragmenting memory across
many databases compromises:

- Efficiency, consistency, security, reliability

Oracle AI Database is converged and ideal for agent memory

- Represent **any kind of memory** via relational, graph, JSON, text etc.
- Perform **any kind of retrieval** using SQL, vector and hybrid search



NEW

Oracle Agent Memory Utility

<https://pypi.org/project/oracleagentmemory/>

- The new agent memory package further simplifies memory development
- Includes APIs required for retrieval, evolution, reflection, etc.
- Compatible with popular memory platforms
- Integrated with popular agent frameworks such as Langgraph, CrewAI, N8N

Standard memory core APIs

Mem0

Letta

DBMS_AGENT_MEMORY



Unified
Agent
Memory



Oracle AI
Database


Example: Long-Term Agent Memory (Facts and Preferences)

Example: Long-term memory


- Collects reusable information during conversation
 - Stored information can be retrieved during other conversations
- ➔ Maintain coherence across sessions

During initial conversation, user preference is collected as memory

The agent recalls the user preference in later conversation to provide a personalized answer



Oct 2025



May 2026



Open Agent Specification

<https://github.com/oracle/agent-spec>

Needed: A standard **agent specification**

AI Agents are becoming the **new interface to enterprise data**


- But today, the eco-system is **fragmented**
- Each agentic framework reinvents data access, reasoning, etc.

Teams are adopting a mix of **open-source and commercial tools** accumulating technical debt

Some standards are emerging, but in different layers:

- **MCP** – defines how **agents access tools**, prompts and contexts
- **A2A** – defines how **agents communicate** with each other

Neither defines the **internal structure** of an agent or **its workflow**



An Industry-First Open Agent Specification

Open Agent Spec is an open-source **framework-agnostic specification** for agents and agentic workflows

- Enables agents to be **portable** across platforms (similar to how **ONNX** provides portability for ML models)
- Supports popular protocols such as MCP and A2A, many ongoing integrations

Standard

Each agent follows a clear **specification**, ensuring **consistent behavior** across frameworks and **increasing trust** in workflows

Modular

Agent Spec's **modular design** makes agents independent building blocks, enabling **fast reuse** and **easy composition**

Portable Across Frameworks

Agent Spec **decouples design from execution**, helping deliver more predictable behavior across frameworks


Agent Spec Example

1. Textual definition of a single agent (JSON/YAML)


Example (JSON)

```
1 {
2   "id": "expert_agent_id",
3   "name": "Adaptive expert agent",
4   "description": null,
5   "metadate": {},
6   "inputs": [
7     {
8       "title": "domain_of_expertise",
9       "type": "string"
10    }
11  ],
12  "outputs": [],
13  "human_in_the_loop": true,
14  "llm_config": {
15    "id": "llama_model_id",
16    "name": "Llama 3.1 8B instruct",
17    "description": null,
18    "metadata": {},
19    "default_generation_parameters": {},
20    "url": "url.of.my.llm.deployment:12345",
21    "model_id": "meta-llama/Meta-Llama-3.1-8B-Instruct",
22    "component_type": "VllmConfig"
23  },
24  "system_prompt": "You are an expert in {{domain_of_expertise}}. Please help t",
25  "tools": [],
26  "component_type": "Agent",
27  "agentspec_version": "26.1.0"
28 }
```

2. Composability Patterns: Agent-in-flow integrates agents into orchestrated workflows. (also, in JSON/YAML)



3. Execution on multiple frameworks, interoperability with various 3rd party tools



Open Agent Spec


Ongoing ecosystem integrations



LangChain

Agent Frameworks

[link](#)



AGNTCY

Agent Discovery and
Exchange Platform


[link](#)



AG-UI and A2UI

Framework-Agnostic Agent
UI tools

[link](#)




ArizeAI

Framework-Agnostic Agent
Observability

[link](#)

- Gaining momentum with agent infrastructure providers; enables framework-agnostic: discovery, deployment, UI interoperability, runtime support, tracing/observability
- Many more upcoming integrations ...



Agent Builders

Tools and frameworks to build, connect,
and operationalize agents

Oracle AI Database Private Agent Factory


A no-code containerized platform to help enterprises rapidly build, test and deploy intelligent data-centric agents and workflows

Trusted platform to build and deploy intelligent agents for enterprise data

- Designed and built for Oracle AI Database
- Developer has full control of workflow
- Data and business logic remain in the database for privacy and security
- Uses Open Agent Spec to define agent workflow


Deep Database Integration: all database features can be consumed in agentic workflows

Deploy anywhere: on-premises, on private cloud, on multi-cloud within your own tenancy



Private Agent Factory from authoring to agent execution

How it works




Example

SupremoBite Multi-agent Assistant

Create a **SupremoBite Assistant** that can process customer refund requests


Example: "The pizza was supposed to be delivered by 6:30PM. It is now 7:10PM, and I want a refund"

Needs a **multi-agent** hierarchical system for different workflow elements




Order Verification Agent

Retrieve the current status, delivery timeline, and any delay records for the specified order




Refund Policy Agent

Reference the latest policy and determine if the order's delay or issue qualifies for a refund




Refund Processing Agent


Initiate the approved refund or credit for the specified order, update all relevant records, and log the transaction



Feedback Agent

Present a feedback survey and escalate if negative sentiment or unresolved issues are detected





Oracle Select AI Agent framework


A simple framework to build, deploy, and manage AI agents in Oracle Database

Developers define and invoke in-database agents using **PL/SQL** or **Python**

- Can also import agent definitions from Private Agent Factory

Core components:

- **Tool** – a software feature or integration that enables an AI agent to perform a specific function
- **Task** – a specific action or unit of work assigned to an AI agent to perform as part of an automated workflow
- **Agent** – an actor with a clearly-defined role that performs assigned tasks as part of an overall workflow
- **Agent Team** – one or more agents that may collaborate to complete tasks within an automated workflow




* roadmap



Agentic Solutions


Domain-specific **patterns and apps**
that accelerate time to value

Agent Factory also comes with pre-built data-centric agents




Knowledge Agent

State of the art built-in RAG agent featuring searching AI vectors, business data, public data, etc.



Data Analysis Agent

Understand schema structure, and extract various data Insights automatically



Deep Research RAG Agent

Decomposes complex questions into **sub-questions** to answer iteratively with **adversarial** reasoning

Coming very soon



Deep Research RAG Agent


An Agent that **researches** ways to improve the response

→ Adaptive Querying

→ Multi-step Reasoning

→ Dynamic Retrieval

→ Self-Reflection



Trusted Answer Search matches end-user natural language questions to previously created reports

Uses AI vector search rather than an LLM to find the best matching report based on report descriptions

LLMs are NOT used to directly answer the user question

- But can be used to build the reports, and to guide the content shown in a report





Private AI Services Container

Private AI models for your private data

Oracle Private AI Services Container keeps data private

Many customers cannot send their enterprise data to third-party AI services


Private AI Services Container offloads expensive AI computation to a local private server

- Vectorizes your data
- Runs open-weight LLMs
- Rerank top results to improve relevance
- Identifies key information in text (NER)
- Runs predictive ML models

Access via a simple, OpenAI-compatible REST API

Deploy **anywhere**

- On-premises, air-gapped, private cloud, or public cloud



Open choice

AI Database supports all the leading AI models and tools



OCI Generative AI



Cohere



OpenAI



Google



Anthropic



Meta



Grok



Hugging Face



Microsoft



Amazon Bedrock



Mistral AI



Jina AI



Qwen



LAION



LangChain



ONNX



LangGraph



Pytorch



Ollama



Haystack



TensorFlow

Customers can choose the AI models and tools they prefer,
and call them using APIs or deploy them as private instances for added security

Additional agentic AI capabilities



Innovate faster with

AI designed for data

- Autonomous AI Database Data Science Agent
- Oracle Graph Studio's graph modeler
- AI Semantic Modeling



Safe and trusted

End data risk

- Globally Distributed Exadata Database on Exascale
- Distributed AI Vector Search
- APEX AI Interactive Reports



Open and universal standards

End data lock-in

- MCP servers
- Choice of App-tier Agent Frameworks
- Deployment Platform flexibility
- GoldenGate 26ai

Whose Prompt is it Anyway?

*The time has come, the Database said,
To talk of many things
Of vectors, graphs, and icebergs deep,
And agents pulling strings.*

*Of prompts and skills and react chains,
Of grounding AI acts
Of why all shiny LLMs
Still need transactional facts.*

*We all praise models night and day,
Their fluency? It flies!
But who has trust? Who serves the truth?
The DB! It never lies!*

*For years, I ran your world unseen
Now, AI steals my glory?
It's when your bank account is cleaned,
It's then, that you'll be sorry!*

*So let us end this foolish song
That databases are done
To keep your data safe and strong
The DB? Second to none.*

Thank you



ORACLE