



## QUICK TAKE

# Oracle AI Database 26ai: *Putting agents & data on the same page*

---

### Trigger

Oracle has introduced a major upgrade to **Oracle AI Database 26ai** spotlighting several core themes: building trust with agentic workflows by treating them as native database processes; extending appeal of Oracle's vector data support, especially to non-Oracle developers; and introducing an open source new declarative framework for defining agentic workflows. In this post, we'll explore how this would impact a typical business use case that is being transformed with agentic AI.

### Agents have lots of moving parts

Throwing agents against the database will radically transform how we work with data. But by their nature for handling multistep processes, agentic workflows carry lots of moving parts. We're going to illustrate how several new features of Oracle AI Database 26ai can keep agents aligned using a common business use case as an example: Why are sales to a specific customer going south?

There are many potential explanations when sales drop. Maybe it's because of competitor pricing, bottlenecks in deliveries, changes in downstream demand, or some disruption caused by external factors such as the regulatory or sociopolitical environments.

Traditionally, this would take a series of disjoint queries, with a person interpreting the answer and then drilling down to eventually pinpoint the cause. Put an agent on it and AI can adaptively iterate the query itself based on the results it retrieves, either with or without human supervision. But there's a potential problem: the models that AI uses to drive agents are not deterministic, so can we trust the results? Let's examine how some of Oracle AI Database 26ai's new features build confidence.

Published April 2026

© dbInsight LLC 2026® | dbInsight.io

## Stop playing games of telephone with the database

*Can you be sure that your agent is in sync with the database? Is the agent working with up-to-date transactional data, and once it gets the data, can the agent reliably pick up where it left off?*

The agent running the sales query learns that in the latest quarter, sales to the customer dropped 10%. It formulates a follow-up query based on the degree of drop. But if in the interim, a new \$500,000 sales order comes in; is the agent's query still valid?

Such disconnects can happen because context windows and transaction state are typically stored separately in the app and database tiers, respectively. Agents are essentially playing games of telephone with the database as they may or may not be working with transactionally current data. Or, if an interim result comes back, is the query reliably reflecting the full context from which it has been iterated?

Oracle AI Database 26ai changes the game by bringing the agent context window and transaction state into a **Unified Agent Memory Core** managed by the database. Agents and databases are now both working on the same (memory) page, ensuring that the agents are asking the right questions of the right data. And by the way, Unified Agent Memory Core eliminates the overhead of moving data to external AI cache.

## Don't let that agent go rogue

*As your agent iterates the next query, can you be sure it won't access data that should be off limits to the end user?*

We're looking into the customer's purchasing patterns, and happen to touch upon sensitive credit status financial data. But if the user submitting the agentic query is part of an affiliate sales team, they may not be authorized to access that data.

Oracle AI Database 26ai settles the issue with **Deep Data Security**, a new feature where the database, not the app tier, is the ultimate arbiter of who can see the data. It requires two tokens, one from the identity and access management gateway and the agent. And based on the token for the user's credentials and the permissions for the agent, the rules established in the data tier will decide access. Case closed.

## Let's go to the best source

*If the trusted, precise information is already there, why not use it?*

A challenge with language models is that they are not deterministic and, of course, can hallucinate if they are not trained on relevant data. Vector search, retrieval augmented

generation (RAG), and semantic tiers have emerged as alternatives to prompts relying on frontier models trained on more diverse (and not always relevant) data sets.

For our use case pinpointing why sales to a specific customer have dropped, we probably want to correlate historic data from internal systems with data representing extrinsic factors, such as changes in customer sentiment or social, political, or economic disruptors.

A pair of new features in Oracle AI Database 26ai can guide agents toward deterministic responses. **Trusted Answer Search** redirects the agent to a RAG vector search for identifying the previously created report that best matches the query. In turn, a similar feature from APEX, **APEX AI Interactive Reports** (an upgrade of APEX Interactive Reports) confines those answers to predefined views or reports. Under the hood, a language model converts questions into filters and aggregations mapping to the corpus of existing APEX reports. Existing APEX customers will automatically be updated to this capability.

## For developers

### Low code lives here

In the low code/no code spirit of APEX, Oracle AI Database 26ai introduces **AI Database Private Agent Factory** that enables business analysts to build, deploy, and scale agents and workflows in a similar drag and drop manner that they have experienced with APEX. In turn, the new **Autonomous AI Vector Database** emphasizes the same ease of use theme enabling developers to quickly build apps with simplified vector APIs, while new native support for vector storage on Apache Iceberg extends the reach of the new Autonomous AI Vector Database to vectors residing in open tables on cloud object storage. In contrast to limited capability vector database offerings, developers using Autonomous AI Vector Database, who need support for more data types and development styles (such as spatial and graph) can simply upgrade to Autonomous AI Database with one click.

### Agent portability

Developers will never settle on a single runtime framework, and so when defining agents, Oracle is introducing a new open source spec, the **Open Agent Specification** that lets them build agents and workflows once and run everywhere. Although unveiled in conjunction with the new Oracle AI Database 26ai release, this spec is database-independent. More importantly, it is runtime framework-agnostic, and that's actually the whole point. The new spec provides a way to describe the components of the agent and how they are orchestrated, then provides an abstraction layer that lets developers use popular runtimes like LangGraph, CrewAI, or Oracle's own AI Private Agent Factory (no code development environment). It's too early to pick a winner here given that Microsoft and Open AI are offering their own alternatives.

## Takeaways

So how does Oracle's approach with aligning agents and data stack up? Where it comes to the foundations of Oracle AI Database 26ai's converged strategy, not just with different data types, but also app tier functions, Oracle's path is unique.

For guiding agents toward deterministic results, Oracle AI Database 26ai takes a more multifaceted approach compared to platforms such as Snowflake that abstract a separate semantics tier. And Oracle has a lot of richness there: besides the new Trusted Answer Search and APEX capabilities, it also provides opportunities to annotate schemas (Select AI Annotations); analytic views (where hierarchies are predefined); and supports property (knowledge) graphs. Just as Oracle database converges data, we'd love to see it converge the richness of its semantics.

Converging functions from the app to the database tier, such as Unified Agent Memory Core and Deep Data Security, are important steps toward building trust with AI. Unified Agent Memory Core eliminates the overhead and uncertainty when syncing context windows with transaction state, thereby eliminating the guesswork (or game of telephone) when agents iterate with data. The closest parallel outside the Oracle environment is Databricks Lakebase, in which the agentic state is maintained in a separate database table in a serverless environment.

For security, Oracle AI Database 26ai eliminates ambiguity by enforcing access to privileges at the database-level, a clear advantage when ensuring that as agents iterate queries, they are not straying into forbidden territory. It complements existing features like Oracle's Private AI Services Container, which keeps data within the Oracle security umbrella when it's passed on to third party models running outside the database (which agents are likely to do).

For Oracle customers, 26ai provides stepping stones, both for low- and pro coders, to start getting their feet wet with agents. Unique features such as Unified Memory Core provide a degree of assurance for getting agents and data, literally on the same page. And for those new to Oracle, there are community and developer editions where they can test drive these new features at no risk.

## Author

Tony Baer, Principal, dbInsight

[tony@dbinsight.io](mailto:tony@dbinsight.io)

Linked In <https://www.linkedin.com/in/dbinsight/>

## About dbInsight

dbInsight LLC® provides an independent view on the database and analytics technology ecosystem. dbInsight publishes independent research, and from our research, distills insights to help data and analytics technology providers understand their competitive positioning and sharpen their message.

Tony Baer, the founder and principal of dbInsight, is a recognized industry expert on data-driven transformation. Baer cohosts with Matt Housley the popular [Its About Data](#) podcast that drew over a million views in its first year. *Analytics* named him as a Top Cloud Influencer for four years in a row. *Analytics Insight* named him one of the [Top 100 Artificial Intelligence and Big Data Influencers](#). His combined expertise in both legacy database technologies and emerging cloud and analytics technologies shapes how technology providers go to market in an industry undergoing significant transformation. A founding member of The Data Gang, Baer is a frequent guest on *theCUBE* and other podcast channels.

dbInsight® is a registered trademark of dbInsight LLC.