

APRIL 2026

Architecting Trusted Agentic AI: How Oracle AI Database Powers Secure, Scalable, and Open AI Applications Optimized for Business Data

Stephen Catanzano, Principal Analyst

Abstract: AI agents have rapidly become the top enterprise AI priority, with organizations expecting them to transform productivity, automate complex workflows, and accelerate decision making across the business. But moving agents from pilot to production requires something most organizations do not yet have: a trusted, unified knowledge base tailored to each AI use case. Oracle AI Database addresses this challenge directly by architecting agentic AI capabilities directly into the database itself, enabling enterprises to build, deploy, and scale AI agents that work with real-time business data while meeting the trust, security, reliability, and governance requirements of production environments.

Agentic AI is now the top enterprise priority

The enterprise AI conversation has shifted decisively. Where organizations once focused on generative AI for content creation and search, attention has moved to agentic AI systems capable of autonomous planning, multi-step reasoning, and taking coordinated actions on behalf of the business. **The appeal is practical: Agents can automate complex workflows, accelerate decision making, and deliver efficiency gains that earlier AI approaches could not.**

80% of organizations view AI agents as their top or a high AI initiative priority.

According to Enterprise Strategy Group (now Omdia) research, 80% of organizations said AI agents are a top or high priority compared to other AI initiatives.¹ That level of consensus reflects how clearly the market

understands the productivity and efficiency potential that agents represent.

What is equally telling is where organizations expect that value to come from. As shown in Figure 1, the leading business drivers for AI agent investment are increasing productivity (39%) and improving and automating processes and workflows (38%), followed by improving decision making speed and accuracy (33%) and enhancing customer experience (31%). These are not experimental or speculative use cases. They sit at the heart of how enterprises operate every day.

¹ Source: Enterprise Strategy Group (now Omdia) Research Report, [AI Agents: The Game-changing Generative AI Use Case](#), August 2025. All research and charts in this Showcase have been taken from this report.

Figure 1. Top AI agent business drivers

Which of the following are your organization’s most important business drivers for the use of AI agents? (Percent of respondents, N=350, three responses accepted)



Source: Omdia

The ambition behind these priorities is striking. Nearly all organizations (91%) agreed that AI agents will improve their efficiency and automation, making it one of the most broadly expected technology outcomes our research has recorded. The opportunity is real and widely recognized. What is less clear is how to close the gap between that ambition and a production deployment that actually delivers.

91% of organizations agreed that AI agents will improve efficiency and automation.

The knowledge problem holding agents back

Despite strong intent, most organizations are discovering that deploying AI agents in production is considerably harder than deploying them in a controlled pilot. The challenge is almost never the AI model. It is the fragmented knowledge infrastructure underneath it.

An AI agent needs more than access to data. It needs access to the right data, in real time, across multiple formats and sources, with the confidence that what it retrieves is accurate, consistent, and governed. It needs to remember context and information from prior interactions. It needs to be prevented from accessing information it should not see. And it needs to do all of this at the scale and speed that business operations demand.

Most enterprise environments were not built for this. Business knowledge is scattered across relational databases, document stores, data lakes, and single-purpose vector databases, each with its own access model, security perimeter, management tools, and operational overhead. Connecting an AI agent to this landscape means building and maintaining complex data pipelines that introduce latency, create consistency and security risks, and multiply potential points of failure.

Our research reinforces this dynamic. While organizations hold high expectations for what agents will deliver, 88% agreed that ongoing human oversight remains essential even as agents automate processes. This reflects practical recognition that agents operating on fragmented, poorly governed knowledge cannot yet be fully trusted to act without checks. The unified knowledge base that agents need to function reliably—an organized, governed, and coherent body of enterprise information tailored to each use case—is something most organizations are still working to establish.

Security and privacy: The foundation of agentic AI

The risks that concern organizations most about AI agents are not abstract. Enterprise Strategy Group (now Omdia) research shows that data privacy (37%) and security vulnerabilities (34%) are the top two most common risks organizations cited related to AI agents, ahead of compliance concerns, a lack of transparency, and loss of human oversight. These are the challenges that stop production deployments in their tracks.

Data privacy (37%) and security vulnerabilities (34%) are the top two most common risks related to AI agents.

When an autonomous agent acts on behalf of the business querying sensitive customer data, initiating transactions, and synthesizing information from across the organization, the question of what it can access and what it can do becomes fundamental. Traditional

security approaches were designed for human users and predictable application queries. AI agents, by contrast, generate unpredictable, wide-ranging queries and can be manipulated through techniques such as prompt injections to attempt access they should never have.

The implication for platform selection is clear. When asked what factors matter most in choosing an AI agent platform, 81% of organizations rated security, compliance, and regulatory requirements as critical or very important, the highest rating of any selection criterion in the research. More specifically, 44% called it outright critical, a level of urgency that no other factor comes close to matching. Organizations that treat security as a deployment checkpoint rather than a design principle are setting themselves up for failure.

Oracle's approach: AI and knowledge engineered together

Oracle takes a fundamentally different approach to the challenge of enterprise agentic AI with its AI Database. Rather than positioning the database as passive storage that AI systems query from the outside, Oracle has architected agentic AI capabilities directly into the database engine. The result is a platform where

building, running, and securing AI agents happens as close to the knowledge as possible, eliminating the fragmentation, latency, and security gaps that come from assembling and connecting separate tools.

Oracle organizes its agentic AI capabilities around three practical priorities: helping organizations build and scale agents faster; minimizing the risks agents introduce; and giving them freedom to work with the models, frameworks, and infrastructure they choose.

Key innovations include:

- **Autonomous AI Vector Database:** Seamlessly integrates real-time enterprise data with AI models to build a knowledge base to support any AI and data use case.
- **AI Database Private Agent Factory:** Provides a no-code visual environment for building, testing, and deploying intelligent agents tailored to specific business needs.
- **Unified Agent Memory Core:** Offers a single, secure memory layer that enables agents to store and retrieve knowledge, spanning semantic vectors, conversation history, structured business records, and graph-based relationships.

“Oracle’s latest innovations empower organizations to rapidly build, deploy, and scale secure agentic AI applications that are suitable for reliable production.”

- Stephen Catanzano, Principal Analyst,
Omdia

Minimizing the risk that agents introduce

Given that security tops the AI agent platform selection criteria by a wide margin, Oracle's approach to data risk deserves close attention. The core principle is straightforward: The only place access control can be truly guaranteed is at the knowledge layer itself, not in application code that agents can bypass.

- **Deep Data Security:** Implements end-user-specific privacy rules directly in the database, ensuring agents can only retrieve what a specific user is permitted to see.
- **Trusted Answer Search:** Uses AI vector similarity search to match questions to pre-validated reports, ensuring deterministic, auditable answers for compliance-critical use cases.
- **Private AI Services Container:** Allows organizations to run AI generation and inference entirely within their own environment, ensuring no data leaves their control.

Open standards and multicloud flexibility

Oracle AI Database eliminates AI data lock-in by supporting open standards and running across multicloud, hybrid, and on-premises environments.

- **Vectors on Ice:** Enables native AI vector search on data stored in open standard Apache Iceberg tables without requiring data movement.
- **MCP Server support:** Allows any MCP-compatible AI agent or application to connect securely to Oracle AI Database.
- **Open Agent Specification:** Provides a framework-agnostic standard for defining agents and workflows, ensuring portability across platforms.

Conclusion

The majority of organizations (80%) have made AI agents a top priority. The expected business outcomes are high, with 91% expecting agents to improve efficiency and automation. Additionally, the leading use cases of productivity, process automation, and faster decision making are exactly the workflows that define competitive advantage.

The answer lies in the foundation. AI agents are only as trustworthy as the knowledge base they draw on, and that knowledge base needs to be unified, governed, and secure enough to be trusted with the processes that matter most. Oracle AI Database provides the architecture to move AI from pilot to production with the security, scalability, and openness that enterprise AI demands.

Omdia recommends that enterprise data and AI leaders evaluate Oracle AI Database as a strategic foundation for their agentic AI programs, with particular focus on Deep Data Security, Unified Agent Memory Core, and the AI Database Private Agent Factory as near-term pathways to moving agents from pilot into production. This approach to deploying agentic AI is architecturally superior to relying on external AI caches and shuffling data from one isolated specialty database to another and sending agents on trips to these multiple databases to retrieve answers with varying degrees of security permissions. A unified data architecture with AI engineered directly into it is the clear winning strategy for the agentic AI era.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.