NUCLEUS
RESEARCH

Oracle AI Database drives 87 percent faster data refresh

ANALYST
Alexander Wurm

The Bottom Line

AI workloads are extending across enterprise infrastructure, and the data primitives that power them, vectors, embeddings, and knowledge graphs, now sit alongside the relational and document data organizations already manage. Organizations running these workloads across fragmented or single-purpose database systems face compounding costs in data latency, operational overhead, and redundant infrastructure that erode both financial agility and competitive positioning. To investigate the value of addressing this challenge, Nucleus interviewed Oracle AI Database customers and found that consolidating transactional, analytical, and integration workloads onto Oracle's converged data platform reduced data refresh times by 87.5 percent by eliminating inter-tool latency, compressed database migration timelines from over a year to two months, and replaced fragmented toolsets with a single engine that natively supports vector search and generative SQL. As vector search, agent memory, and generative capabilities become standard database primitives, organizations that have already consolidated their data platforms are positioned to absorb AI workloads without a second integration cycle.

Overview

Every enterprise AI workload runs on vectors for retrieval, graphs for knowledge, and tokens for generation. These primitives sit alongside the relational and document data that organizations already manage. Running them in separate systems creates the same fragmentation that enterprises spent the last decade consolidating out of their transactional and analytical stacks. Converged database platforms that handle all workload types in a single engine eliminate this second round of fragmentation before it takes root. The question for organizations that have already consolidated is what to build on the unified platform they now operate.

Every enterprise AI workload runs on vectors for retrieval, graphs for knowledge, and tokens for generation

Oracle Database addresses these pressures through a converged architecture spanning relational, vector, graph, JSON, and spatial workloads in a single engine. The platform delivers AI capabilities at no additional charge. Generative SQL through Select AI, native vector search with hybrid retrieval, and agent memory integration with frameworks like LangGraph and mem0 are built into the engine. The autonomous database offering reduces operational burden through self-tuning, self-patching, and self-securing capabilities. The platform deploys across OCI, AWS, and Azure with universal credits across more than 200 regions.

Oracle AI Database

Oracle's converged data platform runs relational transactions, vector search, graph queries, and JSON document operations within a single database instance. Select AI translates natural language queries into SQL, and AI Vector Search enables semantic retrieval alongside structured lookups. The Oracle AI Database 26ai release extends these capabilities with agent memory backed by knowledge graphs, the Open Agent Spec for cross-platform agent portability, and Deep Data Security with row, column, and cell-level access controls for both human users and AI agents. Private AI Services Container enables on-premises vector embedding generation using open weight LLMs.

Select AI translates natural language queries into SQL with vector search enabling semantic retrieval alongside structured lookups

Oracle Autonomous AI Lakehouse extends the platform into analytical and data integration workloads. Managed data ingestion, transformation, cataloging, and delta share capabilities consolidate what previously required a patchwork of ETL tools and custom pipelines. The lakehouse operates as a managed service, reducing the operational burden on database teams and enabling faster time-to-insight for downstream reporting and planning systems.

Benefits

Nucleus interviewed Oracle AI Database customers and found that organizations came to the platform to solve operational problems, consolidating fragmented toolsets, eliminating data latency, recovering from provider failures, and left with infrastructure that natively supports vector search, agent memory, and generative SQL. The operational value and the AI readiness are not separate outcomes. They are the same consolidation.

Consolidating the data pipeline into a single engine reduced refresh times by 87.5 percent, from hours to minutes

Converged Data Architecture

Organizations running separate systems for transactions, analytics, ETL, and emerging AI workloads pay an overhead tax on every data movement between them. Each tool requires its own schema mapping, security model, and operational staff. When the database engine handles relational, analytical, vector, and document workloads in a single instance, that overhead disappears. Organizations that consolidated onto Oracle's converged platform replaced fragmented integration toolsets on the analytical side and migrated entire application portfolios to a single engine on the transactional side. The result is fewer failure points, lower managed service costs, and a platform where vector search and generative SQL are available without provisioning a separate system.

Data Refresh and Decision Velocity

When data passes through a chain of extraction, transformation, and loading tools before reaching a reporting layer, latency compounds at every handoff. Finance and operations teams either wait for the next refresh window or act on stale data. Consolidating that pipeline into a single managed lakehouse engine eliminated the inter-tool latency entirely. Organizations that deployed Oracle Autonomous AI Lakehouse reduced data refresh times by 87.5 percent, from hours to minutes. Finance teams closed periods on current data and shifted from reactive reporting to prescriptive operations management. The same consolidated data layer that accelerated refresh also serves as the foundation for vector search and semantic retrieval across the same datasets.

Organizations solved reliability and latency problems and landed on a platform where vector search and generative SQL are native

Migration Speed and Infrastructure Reliability

Single-vendor or legacy database providers that lack multi-model capabilities force organizations to bolt on additional systems for each new workload type, compounding fragmentation and risk. When a third-party provider fails, the migration path to a converged platform is

faster than moving to another single-purpose engine. Organizations that migrated to Oracle Database compressed timelines from over a year to two months with near-zero downtime. Multicloud deployment on existing cloud infrastructure eliminated the need for a separate cloud migration. The platform these organizations landed on supports relational, vector, graph, and AI workloads natively, meaning the migration that solved a reliability problem also eliminated the need for future bolt-on integrations.

Migration timelines compressed from over a year to two months with near-zero downtime on existing cloud infrastructure

Customer Profiles

Nucleus interviewed two Oracle AI Database customers to understand the operational and strategic benefits of the platform. Both organizations operated in environments where database reliability, data latency, and platform fragmentation created measurable business costs.

Global Digital Entertainment Company

This global entertainment company is unifying its intellectual property archives into an AI-native ecosystem to sell assets targeted across customers, regions, and scheduling operations. The decades of static archives were siloed in various data sources, costly to manage, and inaccessible to stakeholders. The organization ran its database infrastructure on commodity hardware in a campus data center and relied on a third-party cloud database provider for its cloud migration.

The prior production environment experienced three outages within 48 hours, forcing the organization to accelerate its consolidation to Oracle AI Database@AWS, located in an AWS data center and aligns with its existing infrastructure. More than two dozen applications with near-zero downtime and disaster recovery were migrated in two months, instead of the original year-long schedule.

Fewer systems to operate means fewer failure points, lower managed service costs, and a single platform for AI workloads

The company is now using Oracle AI Vector Search, enabling semantic querying of unstructured archival assets, and Select AI, translating natural language queries. This combination of AI-native capabilities in Oracle AI Database 26ai enables non-technical business users to access and interact with real-time insights. The company plans to upgrade its Oracle AI Database@AWS deployment from Exadata Database Service to Autonomous AI Database, eliminating managed service provider costs.

Global Health Company

A global health organization operates 3,000 clinics across 30 countries, employs 7,000 associates, and processes over 35 million annual pet visits under a dozen brands. The organization adopted Oracle Autonomous AI Lakehouse to resolve data refresh latency blocking the financial close process.

Before deployment, the organization refreshed financial data every four to six hours, creating bottlenecks for enterprise performance management across Oracle ERP, Workday, Coupa, and Voyager. Reporting also fed downstream systems including Anaplan, Essbase, Power BI, and OneStream, each requiring separate transformations.

The organization deployed Oracle Autonomous AI Lakehouse as a centrally managed platform for data ingestion, transformation, cataloging, and distribution. Data refresh times dropped from four to six hours to 30 minutes. Finance teams now operate on near-real-time data, enabling prescriptive decision-making rather than reactive reporting. The organization uses delta share extensively for data distribution across downstream systems and cited managed cost savings as an additional benefit of the consolidated platform.

Outlook

The next phase of database consolidation will be driven by autonomous operations. As organizations move from managed infrastructure to self-tuning, self-patching databases, the operational overhead that justified managed service providers and dedicated database teams will continue to shrink. Nucleus expects multicloud deployment to accelerate this shift as organizations demand elastic database capacity on whichever cloud runs their applications, without vendor lock-in or data movement penalties.

The AI layer will determine which platforms survive the consolidation. Databases that treat vector search, agent memory, and generative SQL as native workloads will absorb AI demand without forcing organizations into a second round of integration. Databases that do not will become another point solution in the fragmented stack they were supposed to replace. Deep Data Security for AI agents, open agent interoperability standards, and on-premises embedding generation are early signals of this divide. The convergence of transactional, analytical, and AI workloads into a single engine is no longer a product strategy. It is the operating expectation for enterprise databases.

The convergence of transactional, analytical, and AI workloads into a single engine is the operating expectation for enterprise databases