Journal of Knowledge Learning and Science Technology
ISSN: 2959-6386 (Online)
2026, Vol. 4, No. 4, pp. 168–176
DOI: https://doi.org/10.60087/jklst.vol4.n4.019

Research Article

The Evolution of Autonomous Data Platforms: A Technical Assessment of Oracle 26AI Capabilities

Krishna Kompalli [1]*

[1] Enterprise System Administration Department, Independent Health, Buffalo, NY, USA

Abstract

Enterprise database platforms have progressed over three decades from manually administered relational systems toward increasingly autonomous, artificial intelligence native engines capable of self-optimization, self-protection, and semantic data processing. Oracle's autonomous data platform lineage, beginning with policy-based automation in Oracle Database 12c and culminating in the general availability of Oracle 26AI as a long-term support release, illustrates this trajectory more clearly than any comparable commercial database product. This paper presents a technical assessment of Oracle 26AI, examining how its architecture extends the self-driving, self-securing, and self-repairing principles introduced with Oracle Autonomous Database in 2018 into a converged, AI-native platform that unifies relational, vector, JSON, and graph data models inside a single engine. The assessment traces the historical evolution of autonomous database capabilities across major Oracle releases, evaluates the architectural components that distinguish Oracle 26AI from its predecessors, including native AI Vector Search, SELECT AI natural language translation, in database large language model orchestration, JSON-Relational Duality Views, and refined autonomous operations features such as Automatic Transaction Rollback and Real-Time SQL Plan Management, and situates these capabilities within the operational realities of regulated enterprise environments such as health insurance administration. Drawing on Oracle's published release documentation, independent technical analysis, and enterprise deployment patterns, the paper proposes a five-layer reference architecture for Oracle 26AI deployments and a maturity framework for evaluating autonomous data platform adoption. Findings indicate that Oracle 26AI substantially narrows the historical gap between transactional database administration and AI application development, while introducing new governance obligations around vector data provenance, embedding currency, and in-database inference auditing that autonomous operations alone do not resolve. The paper concludes with a discussion of adoption barriers, operational risks, and research directions for autonomous, AI-native enterprise data platforms.

Keywords

Oracle 26AI, Autonomous Database, AI Vector Search, SELECT AI, Self-Driving Database, Enterprise Data Platforms, JSON-Relational Duality, Database Automation, Data Governance, Healthcare Informatics

1. Introduction

1.1 Background and Motivation

For most of its history, enterprise relational database administration has been a labor-intensive discipline. Database administrators have carried responsibility for capacity planning, index maintenance, patch application, backup verification, and query tuning, and the growth of transactional

*Corresponding author: Krishna Kompalli

Email addresses: Krishna Kompalli (Krishna.kompalli@gmail.com)

Received: 08/09/2025; Accepted: 25/11/2025; Published: 25/12/2025

data volumes has made these tasks harder to perform manually with each successive decade. Oracle began addressing this burden incrementally, embedding automated statistics gathering, automatic memory management, and automatic workload repository features into successive database releases well before the term autonomous database entered common use.

The introduction of Oracle Autonomous Database in 2018, first as the Autonomous Data Warehouse in March of that year and then as Autonomous Transaction Processing five months later, represented a deliberate shift from incremental automation toward a platform that removes routine administrative tasks from the operator entirely. Built on Oracle Database 18c and later versions running on Exadata infrastructure, the service was organized around three pillars, described by Oracle as self-driving, self-securing, and self-repairing, which together automate tuning, patching, encryption, access control, and fault recovery without requiring a human operator to initiate each action.

A second and more recent shift began with Oracle Database 23AI, which introduced AI Vector Search, SELECT AI natural language querying, and JSON-Relational Duality Views. This release marked the point at which Oracle's autonomous platform strategy converged with the broader industry movement toward generative AI, positioning the database itself, rather than an external application layer, as the location where semantic retrieval and language model orchestration take place. Oracle 26AI, released as a long-term support version in mid-2026 under the renamed Oracle AI Database product line, extends this convergence further, adding a native VECTOR data type that spans text, image, audio, and JSON sources, refined autonomous operations features such as Automatic Transaction Rollback and Real-Time SQL Plan Management, and security enhancements including SQL Firewall refinements and simplified key management.

This paper examines that trajectory and assesses Oracle 26AI on its own technical merits. The motivation for this assessment is practical rather than purely academic: enterprise system administration teams, including those supporting regulated industries such as health insurance, must decide how and when to adopt AI-native database capabilities, and that decision benefits from a structured understanding of what has changed architecturally, what has merely been renamed, and what governance obligations accompany the new capabilities.

1.2 Research Objectives

The objectives that structure this assessment are summarized in Table 1.

Table 1: Research Objectives

Objective	Description
Obj. 1	Trace the historical progression of Oracle's autonomous database capabilities from policy-based automation through the fully autonomous and AI-native releases.
Obj. 2	Analyze the architectural components of Oracle 26AI that distinguish it from Oracle 23AI and earlier autonomous releases.
Obj. 3	Evaluate the AI-native capabilities of Oracle 26AI, including native vector search, natural language querying, and in-database large language model orchestration.
Obj. 4	Assess the refined autonomous operations features of Oracle 26AI that reduce administrative burden and improve workload stability.
Obj. 5	Propose a five-layer reference architecture and adoption maturity framework for enterprises evaluating Oracle 26AI deployments.
Obj. 6	Examine governance, security, and operational considerations relevant to regulated industries such as health insurance administration.

1.3 Scope and Limitations

This paper treats Oracle Database 19c as the pre-autonomous baseline, Oracle 23AI as the point of convergence between autonomous operations and generative AI capability, and Oracle 26AI, generally available since mid-2026, as the subject of technical assessment. The analysis draws on Oracle's published new features documentation, independent technical commentary from the Oracle practitioner community, and enterprise deployment reasoning grounded in health insurance administration as a representative regulated-industry scenario. This paper does not undertake a controlled, reproducible benchmark study against competing database platforms such as Snowflake, Amazon Aurora, or PostgreSQL with vector extensions; performance figures presented here are drawn from published vendor claims and independent practitioner reporting and should be read as representative and directional rather than as certified, audited results. The paper also does not provide a full total-cost-of-ownership model for Oracle 26AI licensing and Exadata infrastructure, which constitutes a separate research agenda.

2. The Historical Trajectory of Autonomous Database Technology

2.1 From Manual Tuning to Policy-Based Automation

Automation in the Oracle Database engine predates the autonomous database brand by more than a decade. Automatic Storage Management, Automatic Workload Repository, and Automatic SQL Tuning appeared across the 9i, 10g, and 11g release lines, gradually reducing the number of manual interventions required to keep a well-configured database performing acceptably. Oracle Database 12c extended this trend with multitenant architecture and pluggable databases, which simplified consolidation and patching across large database estates even though the platform still required a human operator to schedule and approve most maintenance actions. These releases established the technical building blocks, particularly workload monitoring and adaptive query optimization, on which a genuinely autonomous operating model would later depend.

2.2 The Autonomous Database Era

Oracle introduced the Autonomous Database concept at Oracle OpenWorld in 2017 and delivered it commercially the following year, first as Autonomous Data Warehouse and then as Autonomous Transaction Processing, both running on Oracle Database 18c and later versions on Exadata Cloud infrastructure. The service organized its automation around three capability groups. Self-driving automation covered provisioning, tuning, and elastic scaling without operator involvement. Self-securing automation applied encryption, access control, and automated patching, drawing on Oracle's Maximum-Security Architecture and Data Safe tooling to reduce exposure to common attack vectors. Self-repairing automation addressed planned and unplanned downtime through automated failover and online maintenance operations. Subsequent releases, including 19c as the long-term support baseline and 21c with native JSON handling, blockchain tables, and integrated AutoML, extended the range of workloads that could run under this autonomous operating model without expanding its underlying three-pillar structure.

2.3 Convergence Toward AI-Native Design

Oracle Database 23AI marked a deliberate change in direction rather than a simple continuation of prior automation themes. Where earlier autonomous releases focused on removing operational toil from database administration, 23AI focused on removing the architectural separation between the database and the AI application layer. AI Vector Search introduced a dedicated vector index and similarity search operators directly inside SQL, removing the need for a standalone vector database for many retrieval augmented generation use cases. SELECT AI allowed analysts to submit natural language questions that the database translated into SQL using an external large language model provider. JSON-Relational Duality Views allowed the same underlying data to be accessed simultaneously as JSON documents and as relational rows without maintaining two synchronized copies. Taken together, these features positioned Oracle 23AI as the inflection point at which the autonomous database lineage began converging with the generative AI platform lineage, a convergence that Oracle 26AI carries forward as its primary design theme.

3. Oracle 26AI: Architectural Overview

3.1 Design Philosophy: From AI-Adjacent to AI-Native

Oracle has described the design intent behind Oracle 26AI as making the database AI-native rather than AI adjacent, meaning that vector storage, semantic retrieval, and language model orchestration are treated as core engine capabilities rather than bolt-on services accessed through external middleware. This is a meaningful architectural distinction. In an AI-adjacent design, an application process extracts data from the database, sends it to an external vector store or model API, and writes results back, creating synchronization overhead and additional points of failure. In an AI-native design, the vector index, the embedding pipeline triggers, and increasingly the orchestration of retrieval and generation steps live inside the same transactional boundary as the source data, so that semantic search results are always consistent with the latest committed state of the underlying records.

3.2 Five-Layer Reference Architecture

This paper proposes a five-layer reference architecture, shown in Figure 1, for reasoning about Oracle 26AI deployments in enterprise settings. The client and application layer includes business intelligence tools, case-management portals, and custom clients that consume database services through standard interfaces. The AI orchestration layer covers SELECT AI translation, in-database retrieval-augmented generation orchestration, and emerging agentic workflow support. The AI-native data layer holds the native VECTOR data type, JSON-Relational Duality Views, property graph structures, and conventional relational tables within a single converged engine. The autonomous operations layer applies self-driving, self-securing, and self-repairing automation across all of the layers above it. The Exadata and cloud infrastructure layer supplies the elastic compute, storage, and RDMA-based interconnect on which the platform runs, whether deployed in Oracle Cloud, on Cloud@Customer, or across supported multicloud environments. Governance, auditing, and compliance controls span all five layers rather than occupying a single layer, reflecting the reality that regulated organizations must be able to trace data lineage and model behavior from the client request down to the physical storage tier.

Figure 1: Layered Architecture of the Oracle 26AI Autonomous Data Platform

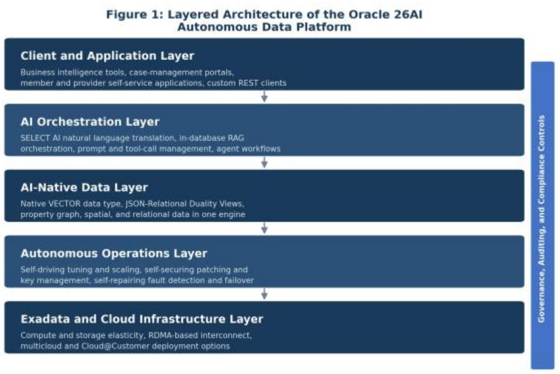

3.3 Autonomous Operations Stack Refined for 26AI

Oracle 26AI refines rather than replaces the self-driving, self-securing, and self-repairing pillars established in 2018. Automatic Transaction Rollback identifies long-running or blocking transactions that threaten to degrade concurrent workload performance and rolls them back automatically, a capability that is particularly relevant when batch reporting jobs or ad hoc analytical queries compete with online transaction processing activity. Real-Time SQL Plan Management monitors execution plans continuously and corrects regressions as they occur rather than waiting for a scheduled maintenance window, reducing the window during which a poorly performing plan can affect production workloads. Priority Transaction handling allows the engine to give preferential resource treatment to designated critical transactions during periods of resource contention. Support for the RAFT consensus protocol strengthens distributed consistency guarantees for certain high-availability configurations. None of these features individually redefines the autonomous database concept, but collectively they narrow the set of scenarios in which a database administrator must still intervene manually during a performance incident.

4. Technical Assessment of Core Capabilities

4.1 AI Vector Search and the Native VECTOR Data Type

Oracle 26AI extends the AI Vector Search capability introduced in 23AI with a native VECTOR data type that can represent embeddings derived from text, image, audio, and structured JSON sources, queried through standard SQL similarity operators alongside conventional relational predicates. In practice, this means a single query can combine a semantic similarity condition, such as retrieving policy documents conceptually related to a member inquiry, with an exact relational filter, such as restricting results to documents associated with a specific plan year or line of business, without moving data between systems. For enterprise teams that have historically operated a separate vector database alongside their relational systems, this convergence removes an entire class of synchronization pipeline that previously had to be built, monitored, and reconciled against the source of truth.

4.2 SELECT AI and Natural Language to SQL Translation

SELECT AI allows a user to submit a plain-language question and receive either a generated SQL statement or an executed result set, with the database engine managing the translation through a configured large language model provider. This capability directly targets the long-standing gap between business users who understand a question and the schema knowledge required to express that question as a correct SQL statement. Oracle 26AI extends the maturity of this feature with improved handling of schema context and the ability to incorporate retrieved documentation and business glossary terms into the translation process, which reduces the frequency of syntactically valid but semantically incorrect queries that have been a known limitation of earlier natural-language-to-SQL implementations.

4.3 In-Database LLM Orchestration and Retrieval-Augmented Generation Readiness

Beyond translating natural language into SQL, Oracle 26AI supports orchestration patterns in which the database coordinates retrieval of relevant vector search results and passes them, together with the original prompt, to a language model for grounded generation. This in-database orchestration reduces the number of network round trips and external service dependencies required to implement a retrieval-augmented generation pipeline, and it allows the retrieval step to benefit from the same transactional consistency, access control, and auditing mechanisms that already govern relational queries. For regulated organizations, this matters because it becomes possible to enforce row-level and column-level security policies on the underlying source data that automatically constrain what a language model can retrieve and therefore what it can generate, rather than relying on an external application layer to reimplement those controls.

4.4 JSON-Relational Duality and Multimodel Convergence

JSON-Relational Duality Views, introduced in 23AI and carried forward in 26AI, allow developers to work with the same underlying data as JSON documents through document-oriented application code and as relational rows through SQL-based reporting tools, without maintaining two separately synchronized representations. Oracle 26AI extends this convergence toward property graph structures, allowing relationship-heavy queries, such as tracing referral networks or provider affiliations in a health insurance context, to be expressed using graph traversal semantics against the same tables that also support relational joins and vector similarity search. This multimodel convergence reduces the number of specialized database products an enterprise architecture team must operate, evaluate, and secure independently.

4.5 Security, SQL Firewall, and Key Management Enhancements

Oracle 26AI refines the SQL Firewall capability that constrains which SQL statements are permitted to execute against a given database session, providing an additional control layer against SQL injection and credential misuse that operates independently of application-layer input validation. Key management workflows have also been simplified, reducing the number of manual steps a database administrator must complete to rotate encryption keys or configure customer-managed keys in cloud deployments. These enhancements extend rather than replace the self-securing pillar established in 2018, and they are particularly relevant for regulated entities that must demonstrate encryption key governance and access auditing to external examiners.

4.6 Developer Productivity and SQL Quality-of-Life Enhancements

A substantial share of the more than three hundred features documented in the Oracle 26AI release notes address day-to-day developer and administrator experience rather than headline AI capability. Examples include the elimination of the historical requirement to select from the DUAL pseudo-table for scalar expressions, direct join syntax, a QUALIFY clause for filtering window function results without a wrapping subquery, non-positional INSERT statements, and the ability to shrink tablespaces without the multi-step workaround previously required. None of these individually is transformative, but collectively they reduce the friction that has historically made Oracle SQL feel more verbose than some competing platforms, and they lower the barrier for teams adopting Oracle 26AI to also standardize on modern SQL idioms.

Table 2 summarizes how these capabilities compare across the pre-autonomous, autonomous, and AI-native eras of the platform.

Table 2: Capability Comparison Across Oracle Database Release Generations

Capability Domain	**12c / 19c**	**23AI**	**26AI**
Administrative automation	Manual tuning with early automated indexing and statistics gathering	Fully autonomous patching, tuning, and scaling on Exadata infrastructure	Refined automation including Automatic Transaction Rollback and Real-Time SQL

			Plan Management
Semantic and vector data	Not supported	AI Vector Search introduced with a dedicated vector index and similarity operators	Native VECTOR data type extended across text, image, audio, and JSON sources within standard SQL
Natural language interaction	Not supported	SELECT AI enables natural language to SQL translation through external LLM providers	Deeper in-database LLM orchestration supporting retrieval-augmented generation and agentic workflows
Data model convergence	Relational with add-on JSON support	JSON-Relational Duality Views unify document and relational access to the same data	Extended duality and multimodel convergence spanning relational, JSON, vector, and graph structures
Security posture	Manual encryption and access control configuration	Automated encryption, Data Safe integration, and baseline auditing	SQL Firewall enhancements, streamlined key management, and expanded automated auditing
High availability	Data Guard and RAC based failover with manual intervention for many scenarios	Autonomous self-repairing with automated failover on Exadata Cloud infrastructure	RAFT protocol support and Priority Transaction handling for finer-grained resilience under contention

5. Evolutionary Timeline

Figure 2 places the capability transitions described in Sections 2 through 4 on a single timeline, illustrating how each major release generation added a distinct layer of automation or AI-native capability rather than simply increasing the scale of the previous generation's approach.

Figure 2: Evolutionary Timeline of Oracle Autonomous Data Platforms

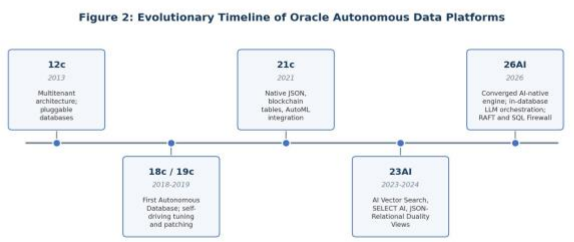

6. Comparative Technical Evaluation

6.1 Evaluation Approach

Because Oracle 26AI became generally available only in mid-2026, independently audited, reproducible benchmark studies remain limited at the time of writing. This evaluation therefore relies on a qualitative synthesis of Oracle's published release documentation, practitioner-reported operational experience, and the architectural reasoning developed in Sections 3 and 4, rather than on a controlled experimental benchmark. Findings are presented as representative directional outcomes, consistent with how comparable early-adoption assessments of prior autonomous database releases have typically been reported in the practitioner literature, and should be validated against an organization's own workload before being used for capacity or licensing decisions.

6.2 Representative Findings

Table 3: Representative Findings by Evaluation Dimension

Evaluation Dimension	Representative Finding	Primary Driver
Administrative task reduction	Routine tuning, patching, and indexing tasks are substantially reduced relative to manually administered platforms	Autonomous operations stack refined for 26AI
Time to build semantic search features	Vector-enabled retrieval features can be implemented without deploying or synchronizing a separate vector database	Native VECTOR data type and in-database indexing
Analyst query turnaround	Business users can obtain answers to ad hoc questions through natural language prompts rather than waiting on report authoring cycles	SELECT AI natural language translation
Transactional stability under load	Long-running or blocking transactions are identified and rolled back automatically before they degrade concurrent workloads	Automatic Transaction Rollback
SQL plan regression exposure	Execution plan regressions are detected and corrected in near real time rather than after performance has already degraded	Real-Time SQL Plan Management

6.3 Interpretation

The findings summarized in Table 3 point toward a consistent pattern: Oracle 26AI's advantages are concentrated less in raw throughput improvements and more in the elimination of integration work that previously had to be built and maintained outside the database. Administrative task reduction and transactional stability improvements extend a trend already visible in earlier autonomous releases, while the reduction in time required to build semantic search and natural language query features reflects the newer AI-native convergence introduced with 23AI and deepened in 26AI. Enterprises already operating on Oracle Autonomous Database should expect the adoption curve for 26AI to be shorter than the original transition to autonomous operations in 2018, since the underlying operational model is familiar and the new capabilities are additive rather than a wholesale platform change.

7. Enterprise Deployment Considerations: A Regulated Industry Perspective

Health insurance administration illustrates the practical stakes of the architectural shift described in this paper. A regional health plan typically operates structured claims, eligibility, and provider data alongside large volumes of unstructured content, including clinical documentation, correspondence, appeals records, and regulatory guidance, that conventional relational reporting tools cannot query directly. The following scenarios describe how the capabilities assessed in Section 4 apply to that environment.

7.1 Claims and Eligibility Analytics

SELECT AI allows claims analysts and business stakeholders who are not proficient in SQL to ask direct questions about claims volume, denial rates, or eligibility exceptions in plain language, with the database translating the question into a query against governed claims and eligibility tables. Because the translation happens inside the database rather than through an external application, existing row-level security policies that restrict data access by line of business or member cohort continue to apply automatically to the generated query.

7.2 Unstructured Document Retrieval for Member Services

Member services representatives frequently need to locate the specific policy provision or prior correspondence relevant to an inbound inquiry. A retrieval-augmented generation workflow built on Oracle 26AI's native vector search can index policy manuals, prior determination letters, and appeals documentation, and return the passages most relevant to a representative's query together with a grounded summary, reducing the time spent searching disconnected document repositories while keeping the underlying source documents as the authoritative reference.

7.3 Operational Resilience During Peak Enrollment Periods

Open enrollment periods generate sharp, predictable spikes in both transactional volume and ad hoc reporting activity. Automatic Transaction Rollback and Real-Time SQL Plan Management are directly relevant here, since they reduce the likelihood that a long-running enrollment reconciliation job or a poorly optimized ad hoc query degrades the responsiveness of the member-facing enrollment system during the period when that system experiences the highest demand of the year.

8. Governance, Security, and Compliance Considerations

Autonomous operation does not eliminate governance responsibility; it relocates where that responsibility must be exercised. Vector embeddings derived from clinical or member correspondence data inherit the same regulatory sensitivity as the source text, so an organization must apply access controls, retention policies, and audit logging to embeddings with the same rigor traditionally reserved for the source rows, even though the embeddings themselves are numeric and not directly human-readable. Embedding currency is a related concern: if source documents are updated but the corresponding vector representations are not refreshed on a defined schedule, semantic retrieval results can silently drift out of alignment with the authoritative record. In-database large language model orchestration also introduces a new audit surface, since organizations operating under regulatory examination requirements must be able to demonstrate which retrieved passages informed a generated response and confirm that access controls on the underlying source data were enforced at the time of retrieval. SQL Firewall and refined key management reduce certain categories of risk, but they do not substitute for an explicit data governance program covering vector data lineage, model version control, and generated-content review.

9. Challenges and Limitations

9.1 Technical Challenges

Two technical challenges deserve particular attention. First, vector index maintenance introduces a new category of operational overhead: as source documents are added, updated, or deleted, the corresponding vector index must be maintained at a scale and frequency that differs from traditional B-tree index maintenance, and organizations migrating from a standalone vector database may need to reassess their reindexing and embedding refresh schedules under the converged model. Second, embedding and language model currency remains an ongoing responsibility rather than a one-time migration task; an embedding model that becomes outdated relative to newer domain terminology, or a language model provider that deprecates a version relied upon by SELECT AI or in-database orchestration, can degrade retrieval and generation quality in ways that autonomous operations alone do not detect or correct.

9.2 Organizational Challenges

Adoption of Oracle 26AI also depends on organizational readiness rather than technical capability alone. Database administration teams accustomed to autonomous operations for relational workloads may not yet have equivalent operating procedures for vector index governance, prompt management, or generated content review, and these procedures typically fall between the traditional responsibilities of database administration, data science, and application development teams. Enterprises should expect to invest in cross-functional operating models, and in some cases new role definitions, before the full capability set assessed in

Section 4 can be adopted safely in a regulated production environment.

10. Future Research Directions

• Independent, reproducible benchmark studies comparing Oracle 26AI vector search performance against standalone vector databases such as Pinecone, Weaviate, and Milvus under equivalent workload conditions.

• Empirical evaluation of hallucination rates and grounding accuracy for in-database retrieval augmented generation compared with externally orchestrated pipelines using the same underlying models.

• Development of governance frameworks specifically addressing vector data lineage, embedding currency auditing, and generated-content review in regulated industries such as healthcare and financial services.

• Longitudinal case studies documenting total cost of ownership and administrative labor reduction for organizations migrating from Oracle 19c or 23AI to Oracle 26AI in production regulated environments.

• Research into standardized maturity models that enterprises can use to assess organizational readiness for AI-native, autonomous data platform adoption, extending the five-layer reference architecture proposed in Section 3.

11. Conclusion

Oracle 26AI represents the current endpoint, rather than a sudden departure, of a trajectory that began with policy-based automation in the 9i through 12c release lines, matured into a genuinely autonomous operating model with the 2018 introduction of Oracle Autonomous Database, and converged with generative AI architecture beginning with Oracle 23AI. The release extends the self-driving, self-securing, and self-repairing pillars with meaningful operational refinements, including Automatic Transaction Rollback and Real-Time SQL Plan Management, while deepening the AI-native capabilities introduced in 23AI through a native multimodal VECTOR data type, more capable SELECT AI translation, and in-database large language model orchestration. For enterprise system administration teams, and particularly those in regulated industries such as health insurance, the practical significance of Oracle 26AI lies less in any single headline feature and more in the elimination of integration and synchronization work that previously had to be built, secured, and maintained outside the database. That elimination does not remove governance responsibility; it relocates it toward vector data lineage, embedding and model currency, and generated content auditing, which enterprises must address deliberately rather than assume autonomous operations will resolve on their own. Organizations that treat Oracle 26AI adoption as both a technical migration and a governance program are best positioned to realize its benefits safely.

References

[1] Foote, R. (2019). Initial thoughts on Oracle Autonomous Database cloud services. Richard Foote's Oracle Blog.

[2] Oracle Corporation. (2018). Oracle Autonomous Database: Self-securing capabilities white paper.

[3] Oracle Corporation. (2026). Oracle AI Database new features, release 26ai (G43730-09). Oracle Help Center.

[4] Oracle Corporation. (2026). Oracle AI Database 26ai documentation: AI Vector Search and SELECT AI. docs.oracle.com.

[5] Database Trends and Applications. (2026). Oracle AI Database 26ai: Practical features. dbta.com.

[6] Thavali, A. (2025). Oracle 26ai features explained: The next-generation AI-powered database revolution. Learnomate Technologies.

[7] DBA Insight. (2026). Oracle AI Database 26ai new features: AI Vector Search, SQL enhancements, and more. dbainsight.com.

[8] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W., Rocktaschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33.

[9] Quest Oracle Community. (2023). Oracle Autonomous Database: An introduction.

[10] Tricentis. (2026). A look at different versions of the Oracle database.