Computer Fraud and Security
ISSN (online): 1873-7056

Oracle AI: Agentic AI and Enterprise Applications Transforming Cloud ERP, Financial Systems, and Compliance Through Autonomous Multi-Agent Architectures

Vijay Tiwari

Senior Manager, Finance Systems Ziff Davis Inc.
ORCID ID: 0009-0009-6584-3539

ABSTRACT

Agentic AI adds another layer of governance risk, compliance obligations and auditability of enterprise financial applications. Oracle Corporation sees agentic AI as the foundation of its enterprise applications strategy. Beginning with embedded AI and task-based agents, we now have outcome-based Fusion Agentic Applications for finance, human resources, supply chain and customer experience. This paper analyzes Oracle's portfolio of agentic AI including: Oracle Fusion Agentic Applications, Oracle AI Agent Studio, Fusion Applications AI Agent Marketplace, OCI Enterprise AI, and Oracle AI Database 26ai from the perspective of financial systems governance, compliance transformation, and regulatory risk. We drew on the literature in multi-agent orchestration, agentic governance, and regulatory technology to assess Oracle's architecture for auditability, role-based access control, deterministic output controls, and data sovereignty in autonomous agent execution pipelines. We conclude that Oracle's full stack SaaS, PaaS, and IaaS integration platform with its large ecosystem of task-specific agents, Fusion Agentic Applications across all major domains of business, and curated marketplace of partner-validated agents provide a meaningful enterprise governance architecture for enterprises involved in regulated financial services, regulated by the EU AI Act, and subject to data sovereignty constraints.

Keywords: Agentic AI, Enterprise Resource Planning, Financial Compliance Transformation, AI Governance, Oracle Fusion Cloud

1. Introduction

Enterprise resource planning systems have been, and remain, the transactional core of most large organisations, capturing, processing and reporting on financial transactions as part of structured and rule-based processes. While these systems are obviously essential for compliance and record-keeping purposes, they are fundamentally reactive: they document what has happened, rather than informing what has to happen next. The current agentic AI can accomplish goal-based thinking, tool usage, and multi-tasking without much human intervention, which makes this model significantly challenged (Abou Ali & Dornaika, 2025; Olujimi et al., 2025).

In a broad set of product releases Oracle has completed a full set of Fusion Agentic Applications for finance, HR, supply chain, and CX. Additional announcements included Oracle AI Agent Studio and an Oracle AI Agent Marketplace of validated partner-built agent applications. With Oracle AI Database 26ai, Oracle Database was positioned as an enterprise AI automation control point. The following sections cover product launches, performance metrics, and deployments. Outside of early adopters, enterprise AI is headed in a platform-first direction: 65.9% of respondents report their intention to follow this strategy. In addition, 38.8% of enterprise buyers expect generative AI to be delivered within agentic systems (Futurum Group, 1H 2026 Survey, 830).

A financial systems architecture, cloud ERP, and an assessment of compliance considerations were used to analyze Oracle's agentic AI strategy. The structure of the paper is as follows: Section 2 presents the theory of agentic AI in the context of enterprise systems. Section 3 discusses the methodology, and section 4 explores Oracle's three-phase journey towards agentic AI. Sections 5-8 critique the agentic products by Oracle, Section 9 outlines governance, compliance, Section 10-12 covers market dynamics, pricing, and the implications of these products for financial leadership.

2. Theoretical Background: Agentic AI in Enterprise Systems

Agentic AI is a subset of AI agents that can autonomously pursue goals, observe the state of the environment, generate plans, and invoke external tools and APIs to execute plans through a chain of actions with minimal human intervention

(Nannini et al., 2026). This contrasts with earlier enterprise AI systems such as predictive analytics engines, robotic process automation, and generative AI copilots. These improve human actions rather than replace the human actor performing those actions.

A systematic literature review guided by PRISMA across 66 studies published from 2019 to 2024 identified three general principles of enterprise agentic AI: autonomy in goal decomposition, coordination across agentic AI agents, and adaptive re-planning based on environmental feedback (Olujimi et al., 2025). Enterprise agentic AI in financial systems operates within multi-entity data environments and involves audit trails and compliance guardrails.

From the perspective of process, the orchestration layer for multi-agent systems (MAS) comprises a planning engine for goal decomposition and sequencing, a policy enforcement module for governing constraints, a state-management function for agent role collaborations and, lastly, a quality operations layer where the outcome of such actions are recorded. Within this orchestrated architecture, the Model Context Protocol (MCP) and the Agent-to-Agent (A2A) protocol are central to orchestrating interoperable, auditable and policy-enforced reasoning across distributed agent architectures, which can be directly mapped to Oracle's AI Agent Studio architecture.

For example, in financial services, agentic crews of specialized agents that handle tasks, including data analysis, model validation, compliance documentation, and audit report writing, can carry out the model risk management in financial services firms while alleviating the administrative burden on human employees (Okpala et al., 2025). Reviewing agentic applications of Artificial intelligence in finance, Aldridge et al. (2026) show progression from trading algorithms to meta-decision systems and continually learning systems with multi-agent coordination. In a PRISMA literature review conducted by Springer for 90 publications from 2018 to 2025, it is evident that neural generative frameworks prevail in finance due to large amounts of data available. Governance model development is considered one of the most crucial gaps in current research (Abou Ali & Dornaika, 2025).

3. Research Methodology

We conducted a systematic mapping study of Oracle's agentic AI ecosystem in the domains of financial systems architecture, cloud ERP interoperability and regulatory transformation, using three modes of evidence to drive our analysis.

The first stream consists of peer-reviewed articles in the Scopus, IEEE Xplore, Springer Nature, Elsevier ScienceDirect and arXiv databases published between 2019 and 2026. Search terms used were "agentic AI", "multi-agent systems", "enterprise resource planning", "AI governance", "RegTech", "financial compliance automation", and "cloud ERP integration". Papers were included based on relevance, soundness, and recency of methods. Preference was given to empirically based studies, systematic reviews, and PRISMA-based research syntheses. Preprints were included if the paper reported the most up to date evidence for new topics given the accepted practice for their rapid dissemination in fast-moving technical areas (Olujimi et al., 2025; Abou Ali & Dornaika, 2025).

The second data stream consists of primary vendor materials provided by Oracle Corporation, such as product announcements and technical architecture specifications and deployment metrics, covering the period from October 2025 through April 2026. These materials were treated as primary source data describing the system under analysis, but not as research evidence in and of themselves.

The third stream was the industry survey data provided by Futurum Group's enterprise AI survey 1H 2026 (n=830) and an independently published November 2025 white paper by the World Economic Forum in collaboration with Capgemini on AI Agent Governance, the latter of which represented an independently gathered quantitative benchmark against which Oracle's position was compared.

We undertook a narrative integration of the three sources of evidence. The Oracle's technology and product claims were matched to theoretical concepts articulated in the peer-reviewed literature. Claims related to governance and compliance were matched to regulations and policies articulated in the peer-reviewed literature: the EU AI Act, GDPR and sector-specific financial regulation. This study did not use any original data collection, research on human subjects, or experiments.

4. Oracle's Three-Stage AI Evolution

Stage 1 - Embedded AI integrated intelligence into applications; surfaced recommendations, flags for anomalies, and predictive insights into customary ERP workflows without changing the essential structure of human-driven decision making.

Stage 2 - AI Agents - workflow driven autonomous systems operate discrete series of transactions across ERP, HCM, SCM and CX Modules that are executed based on the process owner's instructions.

Stage 3 - Fusion Agentic Applications is today's inflection point and a fundamentally new design model for enterprise software. Instead of users clicking through screens and triggering workflows, users will instead state their goals, and coordinate teams of AI agents that reason, decide, and act towards a goal. The most topically-specific, peer-reviewed case study of the impact of combining Oracle AI Agent Studio with Oracle ERP describes this as a basic shift towards agents autonomously identifying anomalies in financial transactions, completing compliance checklists, initiating human resource tasks, and generating documentation in preparation for audits, all while requiring minimal human supervision (Rajagopal Prakashkumar, 2025).

5. Oracle Fusion Agentic Applications by Domain

5.1 Finance and ERP

Oracle's finance agentic applications are designed for the most process and compliance intensive processes in the enterprise. For example, the Claims Settlement Workspace moves exception laden claims processes to smart execution, achieving both cash accuracy and faster settlement cycles. The Collectors Workspace replaces ad-hoc collections management with continuous AI-driven cash flow optimisation which reduces DSO (days sales outstanding). The Cost Accounting Close Workspace surfaces exceptions and next-best actions to shorten the period close process, particularly for CFOs who report quarterly, rather than monthly, figures. The Sourcing Command Center consolidates procurement decisions and supplier negotiations within a single AI-driven interface.

AI automation of select ERP tasks has been shown to be associated with increased predictive accuracy, improved user productivity and lower cost (Mhaskey, 2024). A systematic review of 127 papers on AI and ERP published between 2020 and 2024 concludes that the AI-ERP intersection is moving organizations from legacy automation to clever, collaborative enterprise platforms. This is the architectural vision of Oracle's Finance Agentic Applications (Mannan et al., 2025).

5.2 Supply Chain and Manufacturing

At the Oracle AI World Tour New York on April 9, 2026, the company unveiled 12 new Fusion Agentic Applications in ERP and SCM, including the Warehouse Operations Workspace, which identifies supply chain issues across stock, inbound, outbound and workforce dimensions, and the Contract Compliance Workspace, which performs semantic contract analysis to detect policy violations, prioritize risk and recommend remediation steps, thereby shifting compliance from retrospective audit to proactive risk governance.

LLM-based agentic models can reduce bullwhip effects and outperform centralized demand-based decision-making for inventory management consensus-seeking, thereby delivering independent empirical validation for the operational premise behind Oracle's SCM agentic offering (Jannelli et al., 2025).

5.3 Human Resources and Workforce

The Workforce Operations Agentic Application, announced at Oracle AI World London in March 2026, enables HR leaders to automate manual data collection for payroll and scheduling, accelerate approvals, and eliminate errors in payroll processing. Using transformer-based contextual embeddings, the NLP-enabled ERP platforms have demonstrated an understanding of enough domain-specific business vocabulary to automatically handle HR workflows (Koyeda, 2025).

5.4 Customer Experience

Five Agentic Applications for CX were introduced as part of Oracle's 2023 AI Tech: Oracle AI World Tour New York. They include a Service Request Triage Agent for the routing of service requests to the appropriate agent and an

Escalation Prediction Agent that flags service requests at risk of escalation for an agent's review. Deployment data shows material performance improvement across all CX channels. Quantitative data are in Section 10 metrics table.

Table 1. Oracle Fusion Agentic Applications by Business Domain

Domain	Application	Primary Function
Finance & ERP	Claims Settlement Workspace	Intelligent exception management; improved cash accuracy and reduced settlement cycle time
Finance & ERP	Collectors Workspace	Continuous AI-driven cash flow optimisation; targets reduction in days sales outstanding
Finance & ERP	Cost Accounting Close Workspace	Surfaces material exceptions and next-best actions to accelerate period-end close
Finance & ERP	Sourcing Command Center	Consolidates procurement decisions and supplier negotiations into a single AI-guided interface
SCM	Warehouse Operations Workspace	Proactively surfaces stock, inbound, outbound, and workforce issues with rapid resolution guidance
SCM	Contract Compliance Workspace	Semantic contract analysis for policy deviation detection, risk prioritisation, and remediation
CX	Triage Agent	Intelligent routing of service requests by role and priority
CX	Escalation Prediction Agent	Predictive flagging of high-risk service cases prior to escalation
HR	Workforce Operations Workspace	Reduces manual data gathering; accelerates payroll and scheduling approvals

Source: Oracle Corporation product announcements, March–April 2026.

6. Oracle AI Agent Studio and the Marketplace Ecosystem

Oracle AI Agent Studio is software that enables customers, system integrators such as KPMG, Deloitte, or PwC, and independent software vendors to build, configure, and manage agent-based applications at scale within the enterprise. Users can build agent teams, customize agent actions, and insert human-in-the-loop approval flows for regulatory and compliance-sensitive processes. The Agentic Applications Builder extends this to non-programmers. It allows organizations to launch applications that use AI automation without the programming overhead by using reusable Oracle, partner, and third-party agents. In this way, the Oracle ecosystem has grown to over one thousand agents for different tasks that are available across the Oracle platform (Rajagopal Prakashkumar, 2025).

In October 2025, Oracle released the Oracle Fusion Applications AI Agent Marketplace, a curated marketplace of over 100 validated, partner-designed AI agents from Accenture, Deloitte, IBM, and other partners, that can be discovered and deployed directly inside Oracle Fusion Cloud Applications. Oracle Fusion Applications is the enterprise control plane for

how Oracle-native, third-party and custom agents discover, coordinate and execute cross-functional workflows and tasks.

The standardization of agent communication protocols (MCP and A2A) is an imperative precondition for an interoperable, policy-compliant, multi-agent ecosystem for a marketplace such as Oracle's. Given that 82% of executives expect their organizations will adopt AI agents within three years, the need for this governance and marketplace infrastructure is urgent (Adimulam et al., 2025; World Economic Forum & Capgemini, 2025).

7. OCI Enterprise AI — The Developer Platform for Production Agents

OCI Enterprise AI is an end-to-end platform for developing production ready AI agents at enterprise scale. The agent building process is four steps: choose a foundational model, connect to enterprise data sources and knowledge bases that include structured databases and vector search based retrieval, design tool and API orchestration logic, and deploy into production with IAM-based access control, guardrails, observability, and auditability built in from the first line of code. The platform is built on OCI Superclusters, Oracle's high-density GPU cluster infrastructure, to provide the compute power needed for training, finetuning, and serving large foundation models as part of enterprise workloads without requiring third-party inference service providers.

OCI Enterprise AI has sovereign data hosting and processing options to run enterprise AI workloads, allowing companies in both the highly-regulated financial and government industries to run agentic workloads according to their data residency and jurisdictional compliance needs. Enterprise AI vendors are also required to comply with EU AI Act, GDPR, Cyber Resilience Act, and financial industry regulations. Oracle's sovereign AI architecture can be built to satisfy these compliance requirements (Nannini et al., 2026).

8. Oracle AI Database 26ai — Agentic Data Architecture

Oracle AI Database 26ai (October 2025) and its agentic architecture changes in March 2026 aim for the database to be the enterprise AI automation control center, enabled by three innovations.

8.1 Oracle Unified Memory Core

Oracle Unified Memory Core supplies stateful, persistent memory to agents within the Oracle database engine. Unlike customary agentic architectures, which build context from external vector stores, time lags in the live transactional world mean agents can often be working off stale copies of relevant data. With Unified Memory Core, agents have access to the most current state of the business, including real-time inventory levels, credit availability, and prices. The freshness of the data and architectural coupling of agent reasoning engines with live transactional data layers are both important to the reliability of the financial agents and systemic safety (Gong, 2026).

8.2 Oracle AI Database Private Agent Factory

The Oracle AI Database Private Agent Factory is a no-code agent builder that runs as a containerized service. It is available in a public cloud or on-premises. With this containerized service, users can build three classes of custom agents. Database Knowledge Agents can answer questions related to structured business data. Structured Data Analysis Agents are for analytical reasoning over enterprise data and Deep Data Research Agents are for multi-step research across internal data sources. Private Agent Factory allows users to build and test agents without sending enterprise data to model providers for improved security.

8.3 Oracle Trusted Answer Search

By introducing a deterministically-generative and testable AI answer layer that enables compliance, audit and regulatory supervision, Oracle Trusted Answer Search helps reduce the governance risk enterprises in secure and regulated industries have with probabilistic LLM outputs. This is especially critical in the financial markets, where the lack of deterministic output controls has been identified as a material governance risk (Kurshan et al., 2025).

8.4 Adoption Path and Migration Continuity

Together with Oracle AI Database 26ai, a long-term support release that will replace Oracle Database 23ai, users can apply the release update from October 2025 without needing a database upgrade or application recertification. It also removed a major adoption barrier for enterprise customers in regulated industries: the cost and timeline risk of recertifying deployments.

Table 2. Oracle AI Database 26ai Agentic Architecture Components

Component	Key Capability	Governance Relevance
Unified Memory Core	Stateful, persistent agent memory in-database	Ensures agents operate on current transactional data; reduces stale-context risk
Private Agent Factory	No-code containerised agent builder; on-prem or cloud	Data never transmitted to third-party providers; supports data sovereignty obligations
Trusted Answer Search	Deterministic and testable AI answer layer	Addresses probabilistic LLM reliability concerns in audit and regulatory reporting contexts

Source: Oracle Corporation, October 2025 and March 2026 product announcements.

Governance is not an afterthought in Oracle's agentic approach to business process. It is a core function and role-based access controls, approval workflows, end-to-end process traceability using action logs, observability and measurement of return-on-investment are all built-in features of all Oracle Fusion Agentic Applications.

A central task of Oracle's governance model is the maintenance of the transactional system of record, alongside the agentic execution layer. Compliance with finance, regulation, and securities law requires a deterministic and auditable transactional record. The only prominent difference is in the productivity and workflow execution layer, which uses agent-based self-governance and autonomy to handle failures and unexpected requirements. Other than that, it still complies with the governance layers proposed by Kurshan et al. (2025), enabling agents to act independently without impacting the audit-ready record infrastructure of regulated financial institutions.

To address this governance gap (fewer than 20% of AI agent developers have a formal safety policy and fewer than 10% have a third-party safety review) (Staufer et al., 2026), an enterprise multi-agent governance (TRiSM) framework with an explainability, ModelOps, security and privacy governance model was proposed across the agent lifecycle (Raza et al., 2026). Oracle's architecture follows this framework with explainability through action logs and audit trails, ModelOps through a centralized governance layer in AI Agent Studio, security through IAM-based access control and hosting of sovereign AI data, and privacy through on-premises deployment in the Private Agent Factory.

A regulatory framework for agentic AI in financial markets might include self-regulatory modules, firm-based oversight blocks, regulator-owned monitoring agents, and independent auditing blocks, and would be structurally similar to the Oracle hierarchical governance model (Kurshan et al. 2025). RAG-based agentic compliance systems can automate the precision-based KYC/AML research process and can access relevant regulatory knowledge bases such as the FinCEN guidance and OFAC sanctions lists (Pandey 2025). AI RegTech enables financial institutions to detect fraud, and to monitor transactions with materially improved accuracy (Kothandapani, 2024).

A 12-step compliance architecture for nine enterprise AI agent deployment use cases with different EU AI Act, GDPR and Cyber Resilience Act compliance requirements (Nannini et al., 2026) offers a practitioner-focussed intellectual framework which Oracle's governed deployment model already substantially addresses.

Table 3. Oracle Agentic Governance Architecture Mapped to TRiSM Framework

TRiSM Pillar	Oracle Implementation	Compliance Relevance
Explainability	Step-by-step action logs; end-to-end execution traceability	Supports audit trail obligations under financial regulation and EU AI Act
ModelOps	Centralised governance via AI Agent	Enables consistent policy enforcement across

	Studio	all agent types and workflows
Security	IAM-based access control; sovereign AI data hosting	Addresses GDPR data residency and sector-specific data sovereignty mandates
Privacy	Private Agent Factory on-premises deployment	Prevents enterprise data transmission to third-party model providers

Author's own analysis based on Raza et al. (2026) and Oracle product architecture documentation.

This Oracle push comes just as enterprise buyer behavior is shifting. Futurum Group's 1H 2026 survey of 830 enterprise buyers found that 65.9% of enterprises have a platform-first enterprise AI strategy, 38.8% of enterprise buyers say they expect generative AI to be delivered by agentic systems and 73.7% said they expect to replace their core vendor partners between 2025 and 2028. According to Gartner, agentic AI would answer 80% of routine customer service inquiries by 2029 and deliver 30% in servicing cost savings. 55.2% of enterprise buyers said they would be more confident to allocate more budget to enterprise AI applications if they were more integrated.

Cloud ERP itself yields measurable productivity and profitability gains, the productivity baseline on which the agentic applications of Oracle compound (Sokunbi, 2025). It is only with the arrival of collaborative AI interfacing with ERP that the transition from transaction automation to the next frontier of enterprise intelligence, a new value proposition for cloud ERP systems, begins to take shape (Mannan et al., 2025).

Oracle's differentiated advantages for customers in heavily regulated industries include its full-stack, cloud native SaaS, PaaS, and IaaS portfolio, its data sovereign AI infrastructure, and its pre-integrated compliance controls. In such industries, single-vendor accountability, data sovereignty, and compliance controls tend to trump multi-vendor flexibility.

10.1 Key Market and Deployment Metrics

Table 4. Oracle Agentic AI Ecosystem: Key Market and Deployment Metrics

Metric	Value	Source
Oracle Fusion Agentic Apps available	22+	Oracle, March 2026
New ERP + SCM Agentic Apps (NY World Tour)	12	Oracle, April 9, 2026
CX Agentic Apps launched (NY World Tour)	5	Oracle, April 9, 2026
Task-specific agents in ecosystem	1,000+	Oracle AI Agent Studio
Partner-built agents in Marketplace	100+ validated	Oracle, Oct 2025
Customer inquiries resolved by AI agents	Up to 40%	Oracle deployment data
Improvement in issue resolution time	28%	Oracle deployment data
Increase in first-contact resolution	19%	Oracle deployment data
Platform-first AI adoption (enterprises)	65.9%	Futurum Group, 1H 2026, n=830

Buyers expecting GenAI via agents	38.8%	Futurum Group, 1H 2026
Enterprises considering vendor switching	73.7%	Futurum Group, 1H 2026
Agentic AI CX handling by 2029 (forecast)	80%	Gartner
Operational cost reduction forecast (2029)	30%	Gartner
Better integration = more budget confidence	55.2%	Futurum Group, 1H 2026
Executives planning agent adoption (1–3 yr)	82%	World Economic Forum & Capgemini (2025)

Source: Oracle Corporation; Futurum Group 1H 2026 Survey (n=830); Gartner; World Economic Forum & Capgemini (2025).

11. Customer-Friendly Pricing and Adoption Enablement

Aligned with Oracle's democratizing vision for AI, the embedded AI capabilities and agentic functionality are made available bundled into existing Oracle Fusion application subscriptions at no additional charge to customers. Users of the Oracle Cloud financials, HCM or supply chain and other Fusion applications references receive these free of charge as part of service updates. AI Vector Search in Oracle AI Database 26ai is also included at no additional cost with charges incurred only for additional features such as custom agent development and specialized workflows through Oracle AI Agent Studio.

This pricing model avoids the customary enterprise software friction of separate AI licensing, allowing organizations to begin their agentic journey within their existing investment framework and bypassing the incremental budget approval cycles that have historically inhibited enterprise AI growth.

12. Discussion

The impact on the work of CFOs, financial systems architects, CIOs and management accountants of moving from assisted execution to agentic execution goes beyond productivity impacts. It includes the governance design of the organization and the enterprise risk posture. Four finance agentic application areas are claims settlement, collections, cost accounting close and sourcing. They are the initial generation of enterprise software executing financial workflows. At the business level, agentic AI fundamentally changes the competitive game, e.g., offering early discounts before a human, updating plans on the fly reflecting a cash position, or observing a supply chain disruption without human review and action cycles.

Agentic AI is still in early adopter phase for enterprise. Among executives planning to adopt it, 82% expect to commercially deploy it within the next 1 to 3 years (World Economic Forum & Capgemini, 2025). Further research is needed to understand the implications of full agentic execution, and longitudinal data across a variety of enterprise settings is needed before performance benchmarks can be established.

Governance problems related to coordination failure, prompt-influenced adversarial manipulation, and accountability of delegation from one agent to another, all require formalized trust and risk management strategies (Raza et al., 2026). Likewise, the theoretical underpinnings of the twelve-step compliance framework put forth by Nannini et al. (2026) and the four-layer regulatory framework proposed by Kurshan et al. (2025) must be applied by financial institutions for their Oracle deployments.

Oracle's agentic architecture does have one critical boundary condition. Its governed, full-stack model may favor regulated enterprise customers with Oracle Fusion investments over mid-market enterprises that favor best of breed agent ecosystems that may span multiple vendors. Cross-vendor agent interoperability standards remain an open research and standards question. Financial services firms using hybrid ERP footprints will need to address how Oracle-native agents interact with third-party agent ecosystems before they deploy and scale agentic execution capabilities.

13. Conclusion

Oracle's 1000+ agent-infused Fusion Agentic Applications, covering all major business domains with 22 Fusion Agentic Applications, ranked as the most impactful enterprise software architecture redesign of the ERP era, displacing the previously unchallenged assumption that humans will remain the primary agents for executing enterprise workflows. Their focused AI Agent Marketplace with 100+ partner-validated agents, and agentic AI embedded natively into Oracle AI Database 26ai, give Oracle one of the most complete enterprise agentic stacks in existence.

For financial systems architects, Oracle envisages agentic AI as the new operating model for enterprise finance, where compliance controls, audit logs and regulatory disclosures stay firmly embedded in the core system of record, while the agents optimize and adapt the surrounding processing. There are many interesting angles on this, not least cross-vendor agent interoperability standards, governance of hybrid neuro-symbolic architectures embedded within regulated organizations, and longitudinal productivity impacts of Fusion Agentic Application deployments in a range of financial institutions with varying regulatory difficulties.

This deep embedded governance infrastructure making deterministic output controls, sovereign AI deployment options, and transactional record-keeping possible makes Oracle's platform particularly compelling for regulated financial services institutions and others with high compliance burdens surrounding autonomous agent execution.

Conflict of Interest Statement

The author declares no conflict of interest. This research received no sponsorship, funding, or other support from Oracle Corporation or any affiliated entity. The analysis presented is independent and based solely on publicly available vendor documentation, peer-reviewed literature, and industry survey data.

Funding Statement

This research received no external funding.

References

1. Abou Ali, M., & Dornaika, F. (2025). Agentic AI: A comprehensive survey of architectures, applications, and future directions. Artificial Intelligence Review, Springer Nature. https://doi.org/10.1007/s10462-025-11422-4
2. Adimulam, A., Gupta, R., & Kumar, S. (2025). The orchestration of multi-agent systems: Architectures, protocols, and enterprise adoption. arXiv. https://arxiv.org/abs/2601.13671
3. Aldridge, I., An, J., Burke, R., Cao, M., et al. (2026). Agentic artificial intelligence in finance: A comprehensive survey. arXiv. https://arxiv.org/abs/2604.21672
4. Gong, H. (2026). AI agents in financial markets: Architecture, applications, and systemic implications. arXiv. https://arxiv.org/abs/2603.13942
5. Jannelli, V., Schoepf, S., Bickel, M., Netland, T., & Brintrup, A. (2025). Agentic LLMs in the supply chain: Towards autonomous multi-agent consensus-seeking. International Journal of Production Research. https://doi.org/10.1080/00207543.2025.2604311
6. Kothandapani, H. (2024). Automating financial compliance with AI: A new era in regulatory technology (RegTech). International Journal of Science and Research Archive, 11(1), 2646–2659. https://doi.org/10.30574/ijsra.2024.11.1.0040
7. Koyeda, V. (2025). AI-enhanced ERP systems: Transforming enterprise operations through intelligent integration. Journal of Computer Science and Technology Studies, Al-Kindi Publishers.
8. Kurshan, E., Balch, T., & Byrd, D. (2025). The agentic regulator: Risks for AI in finance and a proposed agent-based framework for governance. arXiv. https://arxiv.org/abs/2512.11933
9. Mannan, M. A., Shaikh, A. K., Khan, A. A., Raahemifar, K., & Mohamed, M. A. (2025). Transforming ERP systems with collaborative AI: Paving the path to strategic growth and sustainability. Computer Standards & Interfaces, Elsevier. https://doi.org/10.1016/j.csi.2025.001444
10. Mhaskey, S. V. (2024). Integration of artificial intelligence (AI) in enterprise resource planning (ERP) systems: Opportunities, challenges, and implications. International Journal of Computer Engineering in Research Trends, 11(12), 1–9.
11. Nannini, L., Smith, A. L., Maggini, M. J., Panai, E., Feliciano, S., Tiulkanov, A., Maran, E., Gealy, J., & Bisconti, P. (2026). AI agents under EU law. arXiv. https://arxiv.org/abs/2604.04604

12. Okpala, I., Golgoon, A., & Kannan, A. R. (2025). Agentic AI systems applied to tasks in financial services: Modeling and model risk management crews. arXiv. https://arxiv.org/abs/2502.05439
13. Olujimi, P. A., Owolabi, P. A., Mogase, R. C., & Wyk, E. V. (2025). Agentic AI frameworks in SMMEs: A systematic literature review of ecosystemic interconnected agents. AI (MDPI), 6(6), Article 123. https://doi.org/10.3390/ai6060123
14. Pandey, V. (2025). Agentic AI with retrieval-augmented generation for automated compliance assistance in finance. International Journal of Science and Research Archive, 15(2), 1620–1631. https://doi.org/10.30574/ijsra.2025.15.2.1522
15. Rajagopal Prakashkumar, P. K. (2025). Analytical study of adapting Oracle AI Agent Studio into Oracle ERP overview. International Journal of Entrepreneurship, Innovation, and Business Strategies, 3(1), 90–103. https://doi.org/10.58425/ijeib.v3i1.452
16. Raza, S., Sapkota, R., Karkee, M., & Emmanuilidis, C. (2026). TRiSM for agentic AI: A review of trust, risk, and security management in LLM-based agentic multi-agent systems. AI Open (ScienceDirect/Elsevier). https://doi.org/10.1016/j.aiopen.2026.S2666651026000069
17. Sokunbi, A. J. (2025). The impact of cloud-based ERP systems on financial performance and operational efficiency. World Journal of Advanced Research and Reviews, 28(3), 1067–1073. https://doi.org/10.30574/wjarr.2025.28.3.1789
18. Staufer, L., Feng, K., Wei, K., Bailey, L., Duan, Y., Yang, M., Ozisik, A. P., Casper, S., & Kolt, N. (2026). The 2025 AI Agent Index: Documenting technical and safety features of deployed agentic AI systems. arXiv. https://arxiv.org/abs/2602.17753
19. World Economic Forum & Capgemini. (2025, November 27). AI agents in action: Foundations for evaluation and governance [White paper]. https://www.weforum.org/publications/ai-agents-in-action-foundations-for-evaluation-and-governance/