International Journal of Advanced Artificial Intelligence Research (IJAAIR)
eISSN: 3087-4270
Volume03 Issue10 | Page No: - 25-36

Received: 23 Sep 2026, Revised: 07 Oct 2026, Accepted: 18 Oct 2026, Published: 30 Oct 2026

DOI: https://doi.org/10.55640/ijaair-v03i10-04

Operationalizing CHAI: Translating Coalition for Health AI Governance Playbooks into Oracle Cloud ERP Controls

Srikanth Dasharadhi
Enterprise Technology Leader, Oracle Cloud ERP and Healthcare Modernization United States

ABSTRACT: Healthcare organizations are enabling AI-embedded ERP features faster than governance rules are being written for them, and in mid-2026 the rules arrived faster than most solution architects expected. The Coalition for Health AI published eight governance playbooks on 27 May 2026, and five days later the Joint Commission opened its voluntary Responsible Use of AI in Healthcare certification, which means a surveyor can now ask a health system to demonstrate its AI governance rather than describe it. Neither document specifies which system setting enforces any of it, and arguably neither should, since both are written to apply across every vendor a health system might run. This paper closes that gap for one platform. Six governance domains running through the CHAI playbooks and RUAIH guidance, namely accountability, transparency, human oversight, pre-deployment risk assessment, continuous monitoring, and third-party risk, are each mapped to a control that already exists inside Oracle Cloud ERP: role-based security profiles, approval hierarchies, and the execution traceability introduced with Oracle's 2026 Fusion Agentic Applications and AI Agent Studio. The mapping is supported by a four-tier risk model that determines how much autonomy an agent may hold, an audit evidence model stating what a surveyor should be shown for each domain, and a six-step implementation sequence. The contribution is a practitioner synthesis drawn from Oracle Cloud ERP transformation programs in regulated health systems. It is not a controlled study, and no compliance outcome, audit finding rate, or certification result is reported or claimed.

KEYWORDS: Coalition for Health AI, Oracle Cloud ERP, Fusion Agentic Applications, AI Agent Studio, responsible AI governance, RUAIH certification, agentic AI, healthcare compliance, HIPAA, HITRUST, NIST AI RMF, AI auditability.

1. Introduction

Healthcare ERP platforms stopped being simple systems of record some time ago. Oracle's 2026 rollout of Fusion Agentic Applications, comprising more than 600 pre-built AI agents spanning Financials, Human Capital Management, Supply Chain, Procurement, and Customer Experience, places autonomous decision-making directly inside the transactional backbone that health systems already use to run claims, procurement, payroll, and revenue cycle work [6], [8], [9]. Named workspaces such as the Cost Accounting Close Workspace, the Sourcing Command Center, and the Warehouse Operations Workspace, together with individual agents including the Ledger Agent and Payables Agent, do not simply surface a recommendation and wait [10]. They execute multi-step processes inside existing approval chains, pausing for a human only where the situation calls for judgment [6].

That shift changes the compliance conversation. An agent that settles a claim, reconciles a vendor payment, or reallocates warehouse stock is making decisions that previously belonged to a person operating under institutional policy. When a person performed that work, hiring standards, training, supervision, and disciplinary process backstopped it implicitly. None of that applies automatically to a software agent inside the ERP. It has to be built.

The Coalition for Health AI, founded in 2021 and now counting well over 3,000 healthcare organizations as members, has been the most active body addressing that gap. With the Joint Commission, CHAI released the Responsible Use of AI in Healthcare guidance in September 2025 [1], [2], and followed it on 27 May 2026 with eight in-depth governance playbooks covering AI

policy, vendor risk, deployment assessment, and monitoring, developed through workshops with more than 150 health AI leaders across in excess of 100 organizations [3]. Five days later, on 1 June 2026, the Joint Commission launched the voluntary RUAIH certification, open to its network of more than 22,000 accredited organizations [12].

That sequence matters more than either document alone. Governance guidance that a health system may read at its own pace is a different proposition from a live certification a surveyor can ask about within the current survey cycle.

What the playbooks and RUAIH do not do, and arguably should not attempt, since they are written to apply across every vendor a health system might run, is identify which Oracle Cloud ERP setting enforces any of it. The gap is not a deficiency in the guidance. It is a translation problem, and translation is a platform-specific exercise.

This paper addresses that translation. The question in the field is not whether to admit AI agents into finance, HR, and supply chain workflows, since the vendor roadmap has settled that, but how to configure them so that a compliance officer, an auditor, or a surveyor sees something governed rather than something running unsupervised. Four contributions follow.

A six-domain mapping from CHAI and RUAIH principles to the specific Oracle Cloud ERP mechanism that enforces each, naming the agent, workspace, or command center where that mechanism lives.

A four-tier risk model determining how much autonomy an agent may hold and which controls become mandatory at each tier, since the mapping alone does not indicate how tightly any given use case should be constrained.

An audit evidence model stating, for each domain, what a surveyor or internal auditor should be shown, which is the artifact the mapping exists to produce.

A six-step implementation sequence, with a worked illustration, for building the mapping into a live program.

Section II positions the work. Section III presents the mapping. Section IV develops the risk model. Section V gives the implementation sequence. Section VI specifies the audit evidence model. Section VII addresses monitoring. Sections VIII and IX cover the worked illustration and the multi-entity case. Sections X to XII discuss, state limits, and conclude.

Note on scope and evidence

What this is. A practitioner synthesis mapping published governance guidance to configurable controls on one ERP platform, drawn from experience delivering Oracle Cloud ERP transformation, data migration, and release management programs in regulated environments.

What it is not. Not a controlled study, not an evaluation, and not a compliance opinion. No claim is made that organizations adopting this mapping experience fewer audit findings, faster certification, or reduced compliance risk, because no such comparison has been performed.

On the case illustration. Section VIII is a composite pattern assembled from recurring engagement experience, not a single documented deployment. It is labeled as such where it appears, and it should be read as a worked example of the method rather than as evidence of outcome.

On currency. Both sides of this mapping are moving. The CHAI playbooks were published in May 2026, RUAIH certification opened in June 2026, and Oracle continues to expand Fusion Agentic Applications on its own release cadence. Section XI states the revalidation cadence this implies.

Not legal or compliance advice. The mapping is a starting point for a health system's governance function to adapt to its own risk posture. It does not replace a formal risk assessment conducted with legal and compliance counsel.

2. Background: the governance gap in healthcare AI adoption

A. Why voluntary frameworks came first

There is still no federal statute setting uniform AI governance requirements for United States health systems. Standard-setting has fallen by default to accreditation bodies, coalitions, and a small number of state legislative attempts.

The Joint Commission and CHAI partnership is the most substantial of these. RUAIH provides a framework that is non-binding but comprehensive for deploying AI safely and ethically [1], [2]. Oklahoma's HB1915, introduced in early 2025, would have required healthcare deployers to operate a formal AI quality-assurance program with continuous performance monitoring and a named governance body. It died in committee, but a state legislature drafting something that specific this early is a

reasonable signal that current voluntary guidance is a preview of future regulation rather than a permanent substitute for it [1].

The broader risk-management vocabulary these frameworks draw on is not new. The NIST AI Risk Management Framework established the govern, map, measure, and manage functions that most sector guidance now echoes [13], and its generative AI profile extends that to systems whose behaviour is less predictable [14]. International management-system standards for AI now exist in parallel [15], [16]. What the healthcare-specific guidance adds is the accreditation hook, and that is what changes the practical urgency.

B. What the playbooks contain

CHAI developed its eight playbooks with input from more than 100 healthcare organizations, spanning academic medical centers, regional systems, and resource-constrained community clinics, so the baseline controls are deliberately scaled to work at different maturity levels [3]. They cover policy-setting, third-party vendor management, deployment risk assessment, and monitoring, and they now feed directly into the RUAIH certification rather than a program still in development [3], [12].

The problem the playbooks address has been described by practitioners as turning good intentions into governed, measurable, sustained practice [12], which is a fair description of the gap a solution architect encounters the moment a policy must become a configuration.

Governance-platform vendors have begun translating the CHAI framework into enforceable controls, with at least one mapping it against HIPAA, HITRUST, and the NIST AI Risk Management Framework [4]. That the market has moved here indicates the same gap is visible commercially. That work sits at the governance-platform layer, however, not inside the ERP where the transaction and the agent both execute.

C. The compressed governance timeline

The interval between guidance and enforceability was unusually short, and it is the reason this translation problem became urgent in 2026 rather than later.

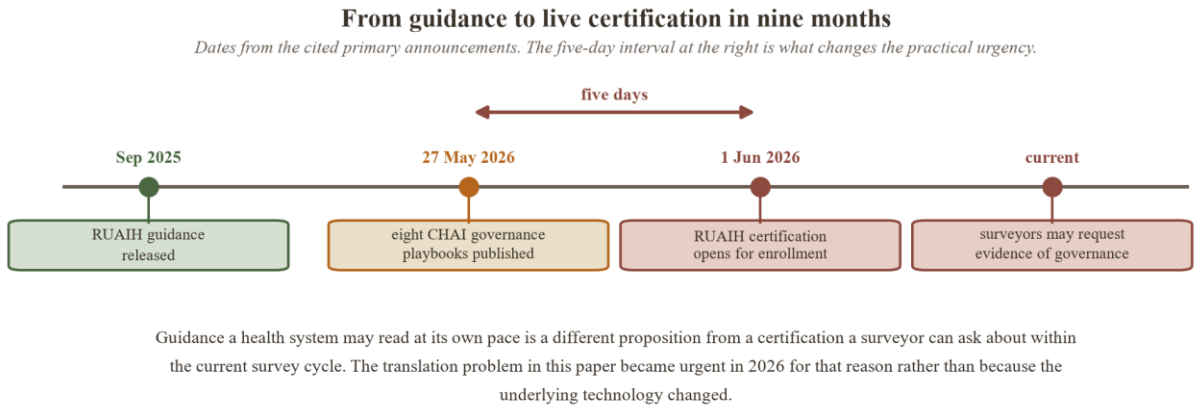


Figure 1. The governance timeline from September 2025 to June 2026. The compression between playbook publication and live certification, five days, is what moves the translation problem from a planning exercise to a survey-cycle obligation. Dates are drawn from the cited primary announcements.

D. Why the ERP is where this must be solved

Most academic literature on healthcare AI governance addresses organizational process: committee structure, escalation paths, and who participates in a risk assessment [5]. That is where governance properly begins. It leaves an open question for any organization running finance, supply chain, or HR on a cloud ERP: once a committee approves a policy, where does that policy take effect?

For an Oracle Cloud ERP estate the answer is increasingly the platform running the transaction, because that is also where the agent making the decision now resides. The auditing literature reached a comparable conclusion about ERP environments well before agentic AI arrived, finding that continuous control assurance in an ERP is a configuration and design discipline rather than a reporting activity layered on afterwards [17], [18].

The enforcement surface does not stop at the ERP boundary. A claims-settlement agent frequently reaches a payer clearinghouse or an electronic health record through Oracle Integration Cloud or a comparable integration platform, so the vendor-risk and traceability controls in Table I must extend across that integration layer rather than terminating at the native ERP screens. Section IX develops this.

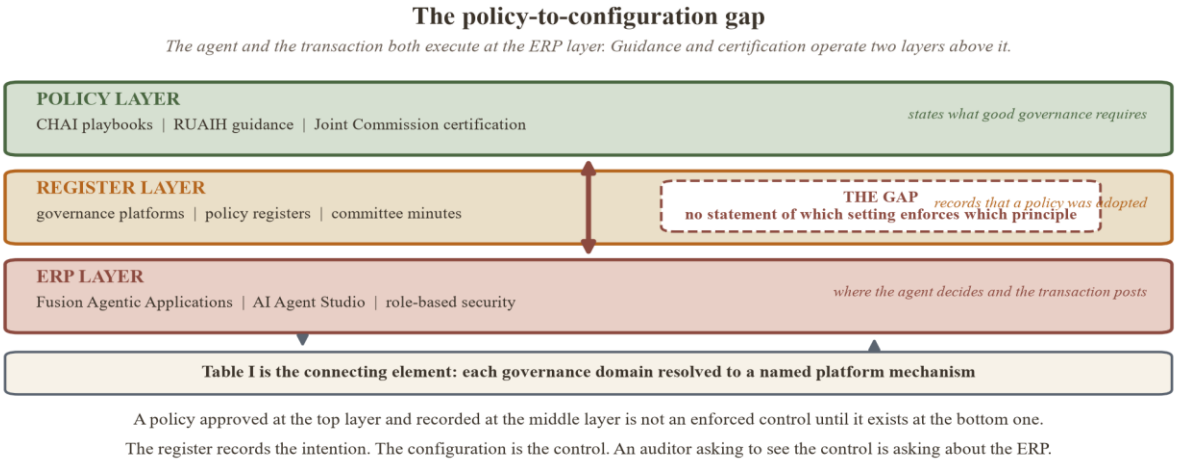


Figure 2. The policy-to-configuration gap. Guidance and certification operate at the policy layer; governance-platform tooling operates at a register layer alongside it; the agent and the transaction both execute at the ERP layer. The mapping in Table I is the connecting element, and its absence is what leaves a policy approved but unenforced.

E. Scope and method

This is a practitioner synthesis drawing on three sources: the published CHAI playbooks and RUAIH guidance [1]-[3], [12]; Oracle's documentation and release announcements describing the security, traceability, and observability built into Fusion Agentic Applications and AI Agent Studio [6]-[11]; and the author's experience delivering Oracle Cloud ERP transformation, data migration, and release management programs in regulated industries.

No primary data was collected, no deployment was instrumented for this paper, and no organization's compliance outcomes were measured. The mapping and the sequence are offered as a starting point for adaptation, not as a validated method.

3. Mapping governance domains to Oracle Cloud ERP controls

Table I is the core of the paper. It narrows the governance domains running through both RUAIH and the CHAI playbooks to six categories, and matches each to the Oracle Cloud ERP mechanism a practitioner would configure, naming the agent, workspace, or command center where that mechanism resides.

Governance domain	CHAI / RUAIH principle	Oracle Cloud ERP control mechanism
Accountability and ownership	Every AI use case requires a named accountable owner and a clear decision authority	Role-based security profiles and approval hierarchies attached to specific agents, including the Ledger Agent, Payables Agent, and Security Command Center, determine who may activate, retune, or approve an agent's

		actions. That role assignment is where accountability is enforced, not in a separate governance register [6], [10]
Transparency and explainability	Decisions must be explainable to affected stakeholders and reviewable after the fact	Execution traceability records every step an agent takes and the reasoning path it followed, so a reviewer reconstructs the decision from the system rather than from recollection or correspondence [6]. The reviewability requirement is well established in the accountability literature [19]
Human oversight and escalation	High-stakes or ambiguous decisions require human judgment before execution	Fusion Agentic Applications advance routine work autonomously and surface only the exceptions and tradeoffs where human judgment changes the outcome [6]. The tolerance band that defines "routine" is a configuration decision, not a vendor default
Pre-deployment risk assessment	Tools require testing and validation for risk before reaching production	AI Agent Studio's testing and validation tooling, together with the prompt playground, are intended to run before an agent is enabled in a live workflow [7]
Continuous monitoring	Deployed AI must be monitored on an ongoing basis, not signed off once	AI Agent Studio observability and the Agent ROI dashboard track process-specific outcomes, such as close duration and exception rates for a Cost Accounting Close Workspace, or RFQ cycle time and negotiated savings for a Sourcing Command Center, rather than a generic usage count [7], [11]
Third-party and vendor risk	External AI vendors and models require the same scrutiny as any other clinical or financial vendor	Because AI Agent Studio supports Oracle, partner, and external agents through a common interface, one governance wrapper, the same roles and the same audit trail, applies

		regardless of who built the underlying agent, including agent logic reached through an integration platform [6], [7]

TABLE I. Governance domain mapping, CHAI and RUAIH to Oracle Cloud ERP. The value of the mapping is that it names the screen or setting a solution architect configures, which allows a compliance officer to ask an auditor's question, show me the control, rather than accept a description of policy.

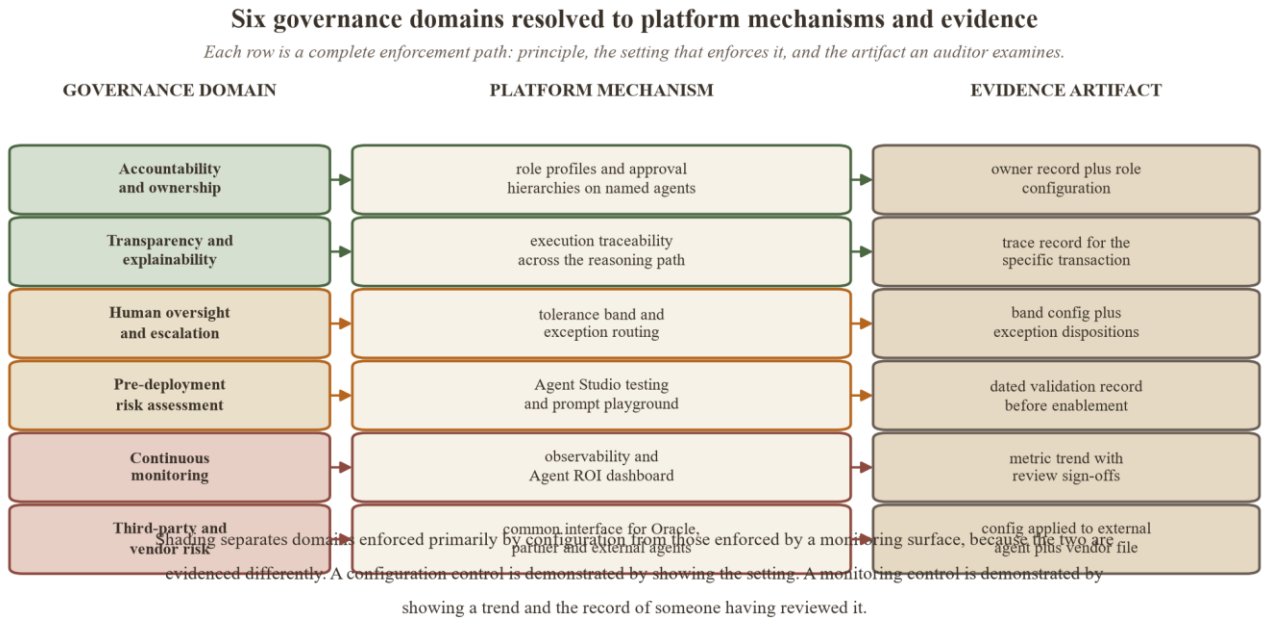


Figure 3. The six-domain mapping rendered as an enforcement path. Each governance principle resolves to a named platform mechanism and produces a specific evidence artifact, which is the object a surveyor examines. Domains that resolve to configuration are distinguished from those that resolve to a monitoring surface, because they are evidenced differently.

4. Scope of the mapping against the bounding question

Table I identifies which control enforces a domain. It does not determine how tightly any individual use case should be constrained, and that is a separate question with its own literature and its own failure modes. How much autonomy an agent may hold, how a tolerance band is set and governed, and how escalation integrity is verified are treated in a companion paper by the author [22] and are not developed here.

What matters for the present paper is only that the two are coupled at one point. The accountability and human-oversight rows of Table I are configured to a depth determined by that risk assessment, so a mapping applied without one produces uniform controls across use cases of very different consequence. The mapping tells a practitioner which setting to configure. It does not tell them where to set it.

5. A six-step implementation sequence

Knowing the mapping does not indicate the order in which to build it, and sequencing matters more than most governance documents acknowledge. The order below reflects what has worked on Oracle Cloud ERP programs in regulated health systems.

Step 1. Build the use case inventory and governance charter. Before any agent is enabled, inventory every Fusion Agentic Application in scope, whether a claims settlement workspace, a collectors workspace, or a Cost Accounting Close Workspace [9], and assign each a named owner. This step is routinely skipped because it resembles paperwork, and it is the one most often regretted.

Step 2. Risk-tier each use case. Apply the model in Table II. The tier determines how tightly the remaining steps are applied, so an error here propagates through everything that follows.

Step 3. Configure role-based access and approval hierarchies. For higher-tier use cases, set security roles so that activation, threshold changes, and exception overrides each require a named approver. This is the accountability and human-oversight rows of Table I made operative.

Step 4. Set audit trail retention and review cadence. Determine how long execution traceability logs are retained and who reviews them, and tie both to RUAIH periodic-review expectations, the organization's survey cycle, and any applicable records obligation under the HIPAA Security Rule [21], rather than selecting an arbitrary interval. Where an agent decision touches protected health information, the retention obligation is set by the stricter of the two.

Step 5. Enable continuous monitoring measured by the process it touches. Observability and the Agent ROI dashboard support a standing review rather than a one-time sign-off [7], [11]. Section VII develops the metric selection, which is where this step usually fails.

Step 6. Run change management deliberately. An agent assuming a process step changes who or what performs the work. Clinical, finance, and HR staff need to understand in plain terms what the agent may decide independently and what it will always return to a person. The clinical AI literature offers a useful precedent here in the model facts label, which communicates a tool's intended use, limits, and validation basis to the people relying on it [20]; the same discipline applied to an ERP agent tells a finance or HR user what the agent is for and where its authority stops.

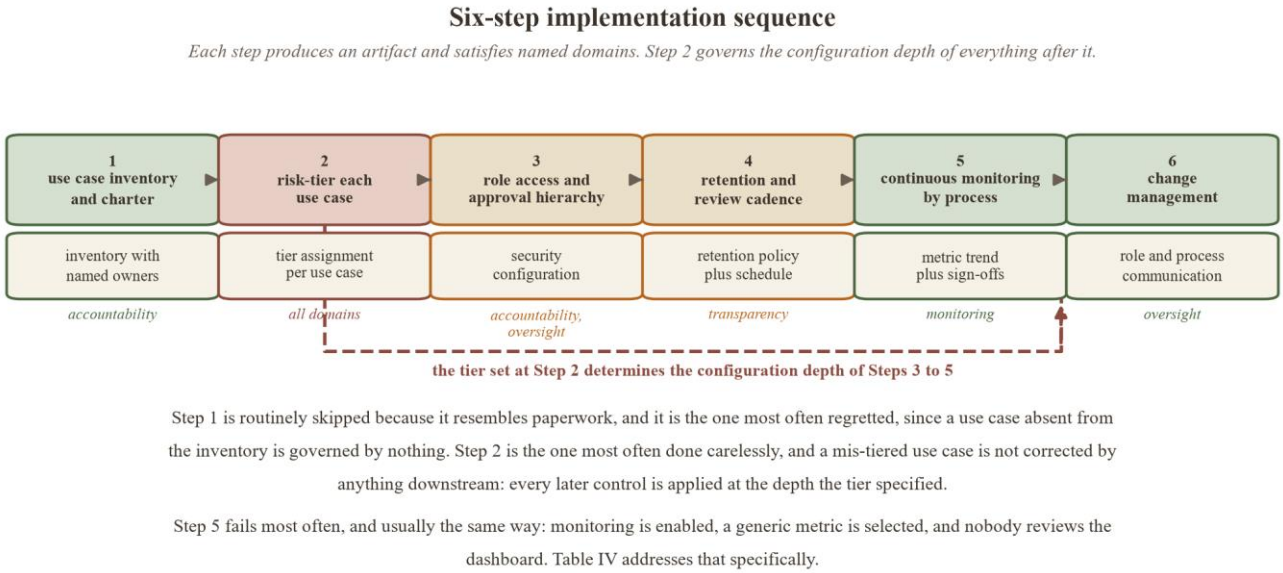


Figure 5. The six-step implementation sequence, with the artifact each step produces and the Table I domains it satisfies. The dependency from Step 2 is drawn explicitly because the risk tier determines the configuration depth of Steps 3 through 5, and a mis-tiered use case is not corrected by anything downstream.

6. The audit evidence model

The mapping and the sequence exist to produce a specific outcome: that when a surveyor or internal auditor asks how AI decisions are governed, the answer is an artifact rather than a description. Stating what that artifact is for each domain converts the mapping from a design aid into an audit-readiness instrument.

Governance domain	What the auditor asks	Evidence artifact	Where it lives
Accountability and	Who is accountable for	Use case inventory entry	Governance charter and

ownership	this agent's decisions, and who may change what it does	with named owner; role assignment and approval hierarchy configuration	platform security configuration
Transparency and explainability	Show me how this specific decision was reached	Execution traceability record for the transaction in question	Platform execution log
Human oversight and escalation	What does this agent decide alone, and what does it escalate	Configured tolerance band and routing rule; exception queue with dispositions	Workspace configuration and exception records
Pre-deployment risk assessment	What testing occurred before this agent was enabled	Validation record from pre-deployment testing, dated before enablement	Agent Studio test artifacts and change record
Continuous monitoring	How do you know it is still performing acceptably	Process-specific metric trend since enablement, with review sign-offs	Observability dashboard and review minutes
Third-party and vendor risk	How is a third-party agent held to the same standard	Role and audit configuration applied to the external agent; vendor assessment record	Platform configuration and vendor risk file

TABLE III. Audit evidence model. The right-hand column is the practical test of the whole approach: if an artifact cannot be located in the platform, the corresponding control is a policy statement rather than an enforced control, regardless of what the governance register asserts.

The distinction this table enforces is the one that separates a governed deployment from a documented one. A policy asserting that high-risk agents require dual approval is not evidence. A configuration showing dual approval attached to the named agent, together with the log of approvals actually granted, is.

7. Monitoring as an evidence source

For the purposes of this paper, monitoring matters because of what it produces rather than what it measures. The continuous monitoring row of Table III requires a metric trend since enablement together with the record of someone having reviewed it, and both halves are necessary. A trend nobody reviewed evidences the existence of a dashboard, not the operation of a control.

Metric selection itself, and specifically why a generic agent usage rate fails to evidence anything a compliance function needs, is treated in the companion paper [22].

8. Case illustration

The following is a composite pattern drawn from recurring engagement experience, not a single documented deployment.

A regional health system moving accounts payable and claims settlement onto Oracle Cloud ERP treated the claims settlement workspace as a single AI use case under

Step 1, with ownership assigned to the revenue cycle director rather than IT. That assignment was deliberate: IT owns the configuration but should not own the risk decision, and conflating the two is a common source of later audit difficulty.

Step 2 placed the workspace at T3 under Table II, since it settles real claims and affects cash accuracy directly [9], and errors reach a payer before a human necessarily sees them.

Step 3 is where the tier became operative. The agent was permitted to settle claims autonomously within a defined tolerance band, with anything outside that band routed to a human collector. This is the same surface-exceptions-only logic Oracle built into Fusion Agentic Applications generally [6], tuned to the workspace's assessed tier rather than left at a default.

Step 4 reused the retention period the organization already applied to claims audits, giving internal audit one familiar cadence to extend rather than a second to learn.

Step 5 tracked the workspace against claims-cycle measures the revenue cycle team already reported monthly, namely days to settle, exception rate, and dollar variance, rather than standing up a separate dashboard that would go unread.

9. Boundaries of the enforcement surface

Two boundaries limit what the mapping in Table I can evidence, and both are developed in the companion paper [22] rather than here.

The first is organizational. In a health system operating several facilities, or an Oracle Cloud instance shared following a merger, risk tolerance is not uniform, so the values behind the controls in Table I vary by entity even where the definitions do not.

The second is technical. A claims-settlement agent frequently reaches a payer clearinghouse or an electronic health record through an integration platform, at which point the decision leaves the ERP's native control surface while remaining part of the same governed action. The third-party risk row of Table I is therefore not satisfied by assessing the agent vendor alone.

Both matter for the present paper in one respect only: an audit artifact that stops at the ERP boundary evidences part of a decision, and the portion outside the boundary is frequently the portion that reaches an external party.

10. Discussion

A. Why ERP-native governance helps

Placing CHAI's principles in the ERP layer rather than a separate governance tool has two practical payoffs. The control resides adjacent to the risk it governs, so the audit team already reviewing financial and procurement controls can extend an established methodology rather than acquire a second toolset. And it narrows the interval between a committee approving a policy and that policy taking effect, since a threshold approved in a given month appears in the live configuration within the same release cycle.

B. What ERP-native governance does not solve

The approach has a boundary worth stating directly, since the paper otherwise argues in its favour throughout.

An ERP-native control governs decisions made inside the ERP. A health system's AI exposure is not confined there. Clinical decision support, ambient documentation, and imaging tools sit outside this surface entirely, and a governance program built only on the mapping in Table I would evidence the back office thoroughly while leaving the clinical estate unaddressed. The mapping is a component of an AI governance program, not a substitute for one.

Nor does configuration substitute for judgment. Table II assigns a tier, but assigning it correctly requires a risk assessment the platform cannot perform.

11. Limitations

General by design, and moving on both sides. The mapping is a starting point for adaptation rather than a specification, and that generality is its principal weakness. The CHAI playbooks were newly published at the time of writing, and RUAIH certification opened for enrollment on 1 June 2026, so there is not yet a track record of how surveyors interpret an ERP configuration during an audit [12]. Oracle continues to expand Fusion Agentic Applications and AI Agent Studio on its own cadence [10], so Table I should be revalidated at every major platform release and at each material revision of the guidance.

Process-level metrics are immature. The measures proposed in Table IV are reasoned rather than validated, and industry practice for measuring agent performance at process level is still developing [11].

No outcome evidence. No claim is made that this approach reduces audit findings, accelerates certification, or lowers compliance risk. Establishing that would require comparing health systems that have implemented an ERP-native mapping against those that have not, which has not been done.

The case is composite. Section VIII is a pattern, not a documented deployment, and should be tested against an organization's own risk assessment rather than copied.

Single platform. The mapping addresses Oracle Cloud ERP. Health systems operating multiple ERPs will need equivalent mappings for each, and Section XII notes this as follow-on work.

12. Conclusion

CHAI and the Joint Commission gave healthcare a governance vocabulary the field trusts, and as of June 2026 they gave it enforcement in the form of a live certification. What was missing was the element connecting that vocabulary to the configuration a solution architect actually applies when standing up Oracle Cloud ERP.

This paper supplies four pieces of that connection: a six-domain mapping to named Oracle Cloud ERP mechanisms, a four-tier risk model determining how much autonomy a use case may hold, an audit evidence model stating what a surveyor should be shown, and a six-step sequence for building it into a program.

The organizing claim is narrow. A policy that cannot be located in a configuration is not an enforced control, whatever the governance register records, and Table III is the paper's instrument for testing that.

Three follow-ons would move this past a mapping exercise. Tracking audit findings and certification outcomes across health systems that have implemented ERP-native governance against those that have not would establish whether the approach affects compliance risk or merely feels more organized. Repeating the mapping for other major ERP platforms pursuing agentic strategies would let multi-ERP health systems apply one governance philosophy rather than several. And now that certification is live, examining whether an actual survey maps cleanly onto these six domains would show whether ERP configuration alone is sufficient or whether a governance-platform layer remains necessary above it.

As Fusion Agentic Applications move healthcare back-office work from assisted to autonomous, organizations treating governance as something configured into the ERP program from the outset will find it substantially easier to demonstrate control than those retrofitting it after an agent is already live.

References

1. Hooper, Lundy and Bookman LLP, "Joint Commission and the Coalition for Health AI (CHAI) Release Governance Frameworks to Address AI in Health Care," Sep. 29, 2025. [Online]. Available: https://hooperlundy.com/joint-commission-and-the-coalition-for-health-ai-chai-release-governance-frameworks-to-address-ai-in-health-care/

2. Joint Commission, "Joint Commission and Coalition for Health AI (CHAI) Release Initial Guidance to Support Responsible AI Adoption Across U.S. Health Systems," Sep. 17, 2025. [Online]. Available: https://www.jointcommission.org/en-us/knowledge-library/news/2025-09-jc-and-chai-release-initial-guidance-to-support-responsible-ai-adoption

3. Coalition for Health AI, "Coalition for Health AI (CHAI) Releases Comprehensive Governance Playbooks to Streamline AI Implementation for Health Systems," PR Newswire, May 27, 2026. [Online]. Available: https://www.prnewswire.com/news-releases/coalition-for-health-ai-chai-releases-comprehensive-governance-playbooks-to-streamline-ai-implementation-for-health-systems-302782394.html

4. Credo AI Advances AI Governance in Healthcare Through Coalition for Health AI (CHAI) Partnership," BusinessWire, Apr. 28, 2026. [Online]. Available: https://www.businesswire.com/news/home/20260428575394/en/Credo-AI-Advances-AI-Governance-in-Healthcare-Through-Coalition-for-Health-AI-CHAI-Partnership

5. S. Freeman, A. Wang, S. Saraf, E. Potts, A. McKimm,

E. Coiera, and F. Magrabi, "Developing an AI Governance Framework for Safe and Responsible AI in Health Care Organizations: Protocol for a Multimethod Study," JMIR Research Protocols, vol. 14, 2025. doi: 10.2196/75702

6. Oracle Corp., "Oracle Introduces Fusion Agentic Applications," Oracle News, Mar. 24, 2026. [Online]. Available: https://www.oracle.com/news/announcement/oracle-introduces-fusion-agentic-applications-2026-03-24/

7. Oracle Corp., "Oracle Expands AI Agent Studio for Fusion Applications with Agentic Applications Builder and New Intelligent Workflow Tools," Oracle News, Mar. 24, 2026. [Online]. Available: https://www.oracle.com/news/announcement/oracle-expands-ai-agent-studio-for-fusion-applications-with-agentic-applications-builder-2026-03-24/

8. Oracle Corp., "Oracle Introduces Fusion Agentic Applications for Finance and Supply Chain," Oracle News, Apr. 9, 2026. [Online]. Available: https://www.oracle.com/news/announcement/oracle-introduces-fusion-agentic-applications-for-finance-and-supply-chain-2026-04-09/

9. Oracle Corp., "Oracle Adds New Fusion Agentic Applications to Help Customers Improve Supply Chain Performance," Oracle News, Jun. 29, 2026. [Online]. Available: https://www.oracle.com/news/announcement/oracle-adds-new-fusion-agentic-applications-to-help-customers-improve-supply-chain-performance-2026-06-29/

10. "26B Roadmaps, Agentic Apps for Finance, HR, SCM, and CX," The Fusion Insider, Oracle Blogs, May 15, 2026. [Online]. Available: https://blogs.oracle.com/fusioninsider/roadmaps

11. "Oracle Brings Governed AI Agents to Fusion Applications," TechWire Asia, Jul. 2026. [Online]. Available: https://techwireasia.com/2026/07/oracle-ai-agent-studio-fusion-applications/

12. Joint Commission, "Joint Commission Releases First of Its Kind, Exclusively Designed for Healthcare Organizations, Voluntary Responsible Use of AI in Healthcare Certification," Jun. 1, 2026. [Online]. Available: https://www.jointcommission.org/en-us/knowledge-library/news/2026-05-responsible-use-of-ai-in-healthcare-certification

13. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1

14. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," NIST AI 600-1, Jul. 2024. doi: 10.6028/NIST.AI.600-1

15. International Organization for Standardization and International Electrotechnical Commission, "ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system," ISO, Geneva, 2023.

16. International Organization for Standardization and International Electrotechnical Commission, "ISO/IEC 23894:2023, Information technology, Artificial intelligence, Guidance on risk management," ISO, Geneva, 2023.

17. J. R. Kuhn and S. G. Sutton, "Continuous Auditing in ERP System Environments: The Current State and Future Directions," Journal of Information Systems, vol. 24, no. 1, pp. 91-112, 2010. doi: 10.2308/jis.2010.24.1.91

18. K. Singh and P. J. Best, "Design and Implementation of Continuous Monitoring and Auditing in SAP Enterprise Resource Planning," International Journal of Auditing, vol. 19, no. 3, 2015. doi: 10.1111/ijau.12051

19. J. Cobbe, M. S. A. Lee, and J. Singh, "Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems," in Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2021, pp. 598-609. doi: 10.1145/3442188.3445921

20. M. P. Sendak, M. Gao, N. Brajer, and S. Balu, "Presenting machine learning model information to clinical end users with model facts labels," npj Digital Medicine, vol. 3, no. 41, 2020. doi: 10.1038/s41746-020-0253-3

21. United States Department of Health and Human Services, "HIPAA Security Rule," 45 C.F.R. Part 164, Subparts A and C.

22. S. Dasharadhi, "Bounding Agent Autonomy in Regulated ERP Environments: A Risk-Tiered Control Model for Healthcare Finance and Supply Chain," companion paper, 2026. [Working title, pending author confirmation.]