

An Agentic ERP Governance Framework for Autonomous AI Agent Deployment in Cloud-Based Industrial Management Systems

Venkata Ramachandra Karthik Chundi

Independent Researcher, Atlanta, GA, USA. ORCID ID: 0009-0008-5238-0590. venkatakarthikchundi@gmail.com

Abstract: Cloud ERP platforms have passed through three distinct automation eras. Scripted batch jobs gave way to robotic process automation, and RPA is now giving way to autonomous agentic AI — systems that reason over enterprise data, select tools dynamically, and execute multi-step business workflows without human direction at every step. The shift is not merely a capability upgrade. Agentic systems behave non-deterministically, invoke tools whose scope may exceed what static governance models anticipate, and can produce cascading process consequences in live financial environments. Governance frameworks built for predictive models and rule-based bots were not designed for this. This paper proposes the Agentic ERP Governance Framework (AEGF), a five-dimension instrument designed to guide the responsible deployment of autonomous AI agents in cloud-based industrial management systems. Drawing on Sociotechnical Systems Theory, the Technology-Organisation-Environment framework, and the NIST AI Risk Management Framework, the AEGF addresses Process Suitability, Autonomy Tiering, Governance and Auditability, Organisational Readiness, and Risk and Continuity Management as an integrated governance architecture. An application to accounts payable automation on Oracle ERP Cloud illustrates how the framework operates in a representative industrial management context.

Keywords: Agentic AI, Cloud ERP, Governance Framework, Autonomous Agents, Oracle ERP Cloud, Digital Transformation, Industrial Management Systems, AI Risk Management

1. Introduction

Automation in enterprise resource planning has never been a single, uniform phenomenon. Different generations of tooling have solved different problems, introduced different failure modes, and demanded different governance responses — and the organisations that have navigated these transitions most successfully have been those that recognised each shift as a governance challenge at least as much as a technical one.

The current shift is more consequential than its predecessors. Agentic AI systems — large language model-based agents equipped with enterprise tool access — do not execute scripts; they reason about tasks, select the tools required to address each step, interpret intermediate outputs, and revise their execution paths when conditions change [1]. The ReAct architecture, which formalises this behaviour as an interleaved sequence of thought, action, and observation, has become the reference design for enterprise agent implementations [1]. Deployed within a cloud ERP platform, a ReAct-pattern agent can retrieve invoice records, cross-reference purchase orders, assess approval thresholds, and initiate payment actions as a coherent, context-aware process rather than a sequence of isolated automated steps. Gartner projects that 40% of enterprise applications will feature task-specific AI agents by 2026, up from less than 5% in 2025 [5]. Oracle's Fusion platform now encompasses more than 600 pre-built agentic applications across its ERP suite, launched in March 2026 [6].

The governance implications of this development have not been adequately addressed in the existing literature. AI governance frameworks developed over the past decade were designed primarily to manage predictive model risk — bias, accuracy, explainability — rather than the dynamic, multi-step, tool-invoking behaviour of autonomous agents



[2], [7]. RPA governance focuses on access control and exception handling for deterministic, rule-bound automation [11], [12]. Neither framework category was built for a system that can reason its way through a novel accounts payable exception at runtime, potentially invoking multiple ERP sub-systems in a sequence that no governance document anticipated. This paper proposes a framework designed specifically for that context.

The contributions of this work are as follows:

- The Agentic ERP Governance Framework (AEGF), a five-dimension structured instrument for evaluating the readiness, risk profile, and oversight requirements of autonomous AI agent deployments within cloud-based industrial management systems.
- An operationalised scoring mechanism across all five dimensions that produces structured deployment tier recommendations, enabling governance decisions to be made on a documented, evidence-based foundation.
- An illustrative application of the AEGF to accounts payable automation on Oracle ERP Cloud, demonstrating the framework's practical utility in a concrete industrial management context.
- A theoretical synthesis grounding the AEGF within Sociotechnical Systems Theory [4], the TOE framework [3], and the NIST AI RMF [2], positioning the contribution within the enterprise AI governance and cloud ERP adoption literatures.

The paper is organised as follows. Section 2 reviews the relevant literature across automation evolution, agentic AI architectures, enterprise AI governance, and technology adoption. Section 3 presents the AEGF. Section 4 applies the framework to an AP automation scenario. Section 5 discusses theoretical and practical implications. Section 6 concludes.

2. Literature Review

2.1 Evolution of Automation in Cloud ERP

Three generations of automation technology have shaped how cloud ERP systems execute business processes, and each has carried its own governance signature. The first — scripted integration and batch processing — produced highly predictable, auditable outcomes, but its brittleness in the face of data variation meant that deviation from expected inputs produced silent failures that required human discovery and remediation.

Robotic process automation addressed some of those brittleness problems by enabling automation at the user interface layer, bypassing the need for formal API integration and extending the automation perimeter to legacy applications. Van der Aalst et al. [12] characterise RPA – software used to simulate interactions between humans and computers to execute tasks according to rules. Siderska et al. [11], in their systematic review, point out the structural restriction in which RPA does not cope with the unstructured input, judgement-based exceptions, and process variability that form a considerable part of ERP workload. Al-Slais and Ali [19] document how these limitations translate into specific governance exposures — audit trail gaps, access control inconsistencies, and failure modes that fall outside the detection scope of conventional IT monitoring. The Oracle Fusion Agentic Applications launch in March 2026 [6], representing more than 600 production-ready agents across the ERP suite, signals that the third generation is no longer a research concept. Bain & Company [9] characterise this shift as a potential inflection point for scaling ERP transformations, contingent on governance maturity matching deployment ambition. Sarferaz [10] examines the technical and organisational preconditions for implementing agentic AI within ERP platforms, identifying integration maturity and data governance as the two most consistently cited adoption barriers.

2.2 Agentic AI Architectures and the ReAct Pattern

What distinguishes agentic AI from prior automation paradigms is not primarily the sophistication of the underlying model; it is the architectural pattern through which the model interacts with enterprise systems. Li et al. [13] define agentic AI as systems characterised by goal-directed behaviour, environmental perception, and

autonomous action selection — a combination that produces qualitatively different operational dynamics from either predictive AI or scripted automation.

The ReAct framework of Yao et al. [1] makes this architecture concrete: by interleaving natural language reasoning traces with action calls to external APIs or tools, the agent maintains contextual coherence across multi-step task sequences. The practical consequence within an ERP context is that an agent does not merely execute a process; it navigates one — reasoning about state, interpreting outputs, and selecting next actions in a way that adapts to the specific conditions of each execution instance. Acharya [16] documents the compounding governance risk that accumulates when organisations deploy agents without addressing traceability requirements, labelling the phenomenon "agent sprawl" and identifying it as an emerging enterprise risk category. Saini's [22] analysis of operating model requirements for agentic AI at scale reaches a convergent conclusion: the governance structures built for predictive AI and RPA are architecturally mismatched with the operational characteristics of agentic systems, and a purpose-built governance layer is required.

2.3 Governance Frameworks for Enterprise AI

The NIST AI Risk Management Framework [2] is the most adopted governance architecture for enterprise AI deployments. Its four functions — Govern, Map, Measure, and Manage — provide a coherent vocabulary for AI governance activities, and its non-prescriptive design allows application across diverse AI system types and organisational contexts. The framework's generality is also its limitation in the agentic context: it does not specify the controls required for non-deterministic, multi-step, tool-invoking agent execution, and it does not address the ERP-specific compliance and audit requirements that make governance gaps in that context materially significant.

Otunlape's [7] examination of AI governance in regulated environments documents this gap empirically, finding that existing frameworks consistently underspecify controls for autonomous multi-step decision systems. The Cloud Security Alliance [15] has extended the NIST AI RMF to address agentic systems specifically, introducing requirements for agent identity management, tool access scoping, and sub-task-level execution logging that the base framework does not cover. Acharya [16] presents a maturity model for governing AI agents in enterprises where auditability, access control, and human control have been found to be the three most important dimensions of governance maturity. Dwivedi et al. [14] provide a foundational multi-disciplinary analysis of enterprise AI adoption challenges, establishing organisational readiness and governance capacity as the two factors most consistently associated with successful outcomes across industries and application contexts.

2.4 Technology Adoption in Cloud Enterprise Systems

The TOE framework of Tornatzky and Fleischer [3] models technology adoption as a joint function of technological characteristics, organisational attributes, and environmental pressures. Its application in the cloud ERP adoption literature is well established. Damali et al. [18] find that technological compatibility and top management support are the strongest adoption predictors in cloud ERP deployments in healthcare environments. Hahoang [17] demonstrates that competitive environmental pressure moderates the relationship between technological readiness and adoption outcomes in SME digital transformation contexts, indicating that adoption decisions are never purely technical assessments.

The relevance of TOE to the agentic governance context extends beyond adoption prediction. Governance framework adoption — the decision to implement a structured assessment instrument like the AEGF rather than deploying agents ad hoc — is itself a technology adoption event governed by the same organisational and environmental factors that shape ERP adoption decisions. Embedding the TOE's organisational dimension as an explicit, scored governance variable (AEGF Dimension 4) rather than treating it as implementation background directly addresses the gap between technical governance design and organisational governance capacity that Amanollahnejad et al. [21] identify as the primary determinant of enterprise AI deployment outcomes.

2.5 Research Gaps

The picture that emerges from this review is one of productive partial coverage. Agentic AI architectures are well characterised at the technical level [1], [13]; enterprise AI governance has a robust foundational literature [2], [7], [15], [16]; and cloud ERP technology adoption is extensively studied [3], [17], [18]. The specific combination — a governance instrument designed for autonomous agent deployment within cloud ERP systems, integrating architectural, governance, and adoption-readiness considerations into a single scored framework — is not represented in the literature. The Agentic ERP Governance Framework proposed in this paper addresses this combination directly.

3. The Agentic ERP Governance Framework (AEGF)

3.1 Design Rationale and Theoretical Foundations

Three bodies of theory converge in the AEGF's design, and understanding their distinct contributions clarifies why the framework takes the form it does. Sociotechnical Systems Theory [4], originating with Bostrom and Heinen's analysis of MIS failures, holds that information systems succeed or fail as joint human-technical systems — that technical capability without commensurate organisational capacity produces governance gaps regardless of the quality of the technical design. The TOE framework [3] establishes that adoption and governance outcomes are functions of technology fit, organisational readiness, and environmental context, not of technical specification alone. The NIST AI RMF [2] provides a governance activity architecture — Govern, Map, Measure, Manage — that the AEGF operationalises for the agentic ERP deployment context.

The five dimensions are not independent assessment criteria. They form an interdependent governance architecture: a high Autonomy Tiering score (Dimension 2) is only actionable if Governance and Auditability (Dimension 3) confirms that monitoring and rollback infrastructure is in place; an aggressive deployment timeline is only defensible if Organisational Readiness (Dimension 4) confirms the governance capacity to manage the deployment. Table 1 maps each dimension to its primary governance concern and theoretical grounding.

Table 1: AEGF Dimension Overview

Dimension	Primary Governance Concern	Theoretical Grounding
1. Process Suitability	Process fit for agent execution	TOE — Technology
2. Autonomy Tiering	Appropriate operational boundary	NIST AI RMF — Map/Govern
3. Governance & Auditability	Traceability and reversibility	NIST AI RMF — Measure/Manage
4. Organisational Readiness	Human and institutional capacity	TOE — Organisation; SST
5. Risk & Continuity	Failure modes and recovery	NIST AI RMF — Manage; SST

3.2 Dimension 1: Process Suitability Assessment

Process Suitability evaluates whether a candidate business process is structurally appropriate for autonomous agent execution. Four sub-criteria are assessed: process structure (predictability of inputs, outputs, and execution path); exception frequency (proportion of instances requiring non-standard handling); financial consequence (monetary exposure of an incorrectly executed instance); and regulatory sensitivity (extent of compliance obligations demanding explicit human accountability).

A high-structure, low-exception, low-consequence process — three-way match invoice processing for invoices below a defined value threshold, for example — is a strong candidate for Tier 2 autonomous execution. A low-structure, high-exception process — vendor dispute resolution for novel contract interpretations, for example — is a Tier 0 or human-execution candidate regardless of the agent's technical capability. Between these extremes, most real-world ERP processes produce a differentiated suitability profile in which specific sub-processes within a broader workflow receive different tier assignments, and the governance design must reflect that granularity.

3.3 Dimension 2: Autonomy Tiering

Autonomy Tiering defines the operational boundary within which an agent executes without human intervention. The AEGF establishes four tiers, drawing on the autonomy spectrum frameworks in [13] and [22]:

Tier 0 — Supervised Execution: The agent creates a plan of actions for approval by humans prior to executing them. None of ERP transactions is performed without human authorization. For use in innovative processes, critical decision-making, and start-up periods.

Tier 1 — Co-Pilot Execution: The agent executes low-risk actions autonomously, while making decisions that go beyond certain financial or compliance limits are routed to humans for approval. Agent's execution is real-time but limited by configuration parameters set during deployment.

Tier 2 — Monitored Autonomous Execution: Governed Autonomous Execution: The agent executes a process autonomously in accordance with certain parameters, with real-time governance monitoring which can override or stop execution once anomalies go beyond specified threshold levels. Humans review exceptions and pre-scheduled periodic reviews.

Tier 3 — Full Autonomous Execution: Fully Autonomous Execution: To be used for processes with greatest structural regularity, lowest risk, maximum reversibility, and best monitoring instrumentation. Governance is periodic and audit-oriented.

Tier assignment is a dynamic governance parameter, reviewed at defined intervals and advanced incrementally as operational evidence accumulates. No deployment begins above the tier recommended by the combined AEGF assessment.

3.4 Dimension 3: Governance and Auditability

Four capabilities are required for a deployment to score adequately on this dimension. Execution trace logging must capture the agent's reasoning steps, tool invocations, input parameters, and output states at the sub-task level — not merely the resulting ERP transactions. In an Oracle ERP Cloud context, this requires integration between the agent orchestration layer (OIC-hosted agent runtime or Oracle Digital Assistant) and the platform's native audit trail infrastructure. Anomaly detection must identify deviations from baseline execution behaviour — unusual transaction volumes, out-of-hours execution, tool invocations outside the defined process scope — and trigger alerts or automatic execution halts. Rollback capability must be tested and documented for each process tier above Tier 0, covering journal entry reversals, workflow cancellations, and payment holds where applicable. Role-bounded access requires that agent service accounts are provisioned at minimum-privilege scope and subject to the same periodic access review applied to human user accounts.

3.5 Dimension 4: Organisational Readiness

Organisational Readiness is the dimension where most agentic ERP deployments are currently most vulnerable. The technical capability to deploy an Oracle Fusion Payables Agent is available to any Oracle Cloud customer as of March 2026 [6]; the organisational capacity to govern that deployment responsibly is considerably rarer. Amanollahnejad et al. [21] make this asymmetry precise: sociotechnical alignment — the match between technical system design and the human and institutional structures operating alongside it — is the primary predictor of enterprise AI deployment outcomes, and its absence is the most common source of deployment failures that technical quality alone cannot explain.

Four sub-criteria capture the relevant organisational factors: AI literacy among ERP process owners, internal auditors, and senior leadership; governance capacity in the form of defined roles, policies, and escalation procedures for AI system oversight; change management maturity as evidenced by track record and current capability; and vendor partnership maturity, which at this stage of the Oracle Fusion Agentic Applications platform's development reflects the quality of the joint problem-solving relationship between the organisation and its Oracle implementation partner [20].

3.6 Dimension 5: Risk and Continuity Management

Three risk categories require explicit assessment and documented mitigation plans. Execution risk — the probability and consequence of an agent executing an incorrect action — is mitigated through the threshold parameters and validation checks embedded in Dimension 2's tier configuration, together with transaction value limits and pre-execution ERP master data validation. Dependency risk — the exposure introduced by the agent's reliance on language model inference services, API connections, and external data sources — is mitigated through graceful degradation procedures that revert affected processes to manual or RPA execution when agent infrastructure is unavailable or degraded. Auditability risk — the risk that insufficient logging prevents post-hoc investigation of agent-executed transactions — is jointly specified with Dimension 3 at the framework design stage, ensuring that monitoring infrastructure is designed as a unified capability rather than assembled from separate controls.

3.7 Mathematical Formalisation of the AEGF

To enhance rigour and provide a quantitatively grounded basis for deployment recommendations, this section formalises the AEGF scoring model using a weighted composite index with dimension interdependency constraints.

3.7.1 Dimension Scoring Model

Let D_i denote the normalised score for dimension i , where $i \in \{1, 2, 3, 4, 5\}$ corresponding to **Process Suitability**, **Autonomy Tiering**, **Governance and Auditability**, **Organisational Readiness**, and **Risk and Continuity**, respectively. Each dimension score is computed from its sub-criteria as:

$$D_i = (1/n_i) \cdot \sum s_{ij}, \text{ where } s_{ij} \in \{1, 2, 3, 4, 5\}$$

where n_i is the number of sub-criteria in dimension i , and s_{ij} is the sub-criterion score assigned during assessment. Each sub-criterion is evaluated on an integer scale from 1 (lowest capability or suitability) to 5 (highest).

3.7.2 Weighted Composite AEGF Score

The composite AEGF score S is computed as a weighted sum of dimension scores:

$$S = \sum w_i \cdot D_i, \text{ where } \sum w_i = 5 \text{ and } w_i \in \mathbb{R}^+$$

In the baseline uniform-weight configuration, $w_i = 1$ for all i , yielding $S \in [5, 25]$ — consistent with the scoring scale in Table 2. Practitioners operating in high-regulatory-sensitivity environments may apply domain-adjusted weights, provided the constraint $\sum w_i = 5$ is preserved and weight assignments are documented and justified in the governance record.

3.7.3 Interdependency Gate Function

A critical architectural feature of the AEGF is that no dimension score is sufficient in isolation. The framework enforces two interdependency constraints through a gate function G :

$$G(D_2, D_3) = 1 \text{ iff } D_3 \geq D_2 - 0.5, \text{ else } G = 0$$

$$G(D_1, D_4) = 1 \text{ iff } D_4 \geq \min(D_1, 3.0), \text{ else } G = 0$$

The first gate (**Autonomy-Auditability Gate**) enforces that an agent's autonomy boundary may not exceed its monitored governance capability by more than 0.5 score points. An organisation with $D_2 = 5$ (full autonomy intent) but $D_3 = 3$ (partial audit infrastructure) triggers $G = 0$, which blocks advancement to Tier 2 regardless of composite score S . The second gate (**Suitability-Readiness Gate**) enforces that organisational readiness must meet a

minimum threshold commensurate with process suitability: if $D_1 \geq 3$ (moderate-to-high suitability), D_4 must also reach at least 3 to avoid a deployment block.

The effective deployment score S^* is therefore defined as:

$$S = S \cdot G(D_2, D_3) \cdot G(D_1, D_4)^*$$

When either gate evaluates to 0, $S^* = 0$ and deployment is blocked at the recommended tier, irrespective of composite score magnitude. This ensures that high performance in unconstrained dimensions cannot compensate for structural governance deficits in the constrained pairs.

3.7.4 Tier Assignment Function

The deployment tier T is a step function of S^* :

$T(S^*) =$

- Tier 3 if $S^* \geq 23$
- Tier 2 if $21 \leq S^* < 23$
- Tier 1 if $16 \leq S^* < 21$
- Tier 0 if $11 \leq S^* < 16$
- No Deploy if $S^* < 11$ or $G = 0$

3.7.5 Tier Progression Criterion

Following initial deployment, tier advancement is permissible when operational evidence accumulated over an observation period Δt satisfies:

$$\Delta T = +1 \text{ iff } (S_{\text{new}} > S_{\text{current}}) \wedge (D_3_{\text{new}} \geq D_3_{\text{current}}) \wedge (E_{\text{anomaly}}(\Delta t) \leq \epsilon)$$

where $E_{\text{anomaly}}(\Delta t)$ is the empirical anomaly rate observed during the review period, and ϵ is the organisation-defined acceptable anomaly threshold (recommended default $\epsilon = 0.02$, representing no more than 2 anomalous executions per 100 agent-driven transactions). This criterion formalises the requirement that tier progression be evidence-based and that governance monitoring capability (D_3) must not regress between review cycles.

3.8 AEGF Scoring and Deployment Recommendation

Each dimension is assessed on a five-point scale across its sub-criteria, producing a per-dimension score and an aggregated AEGF total. Table 2 maps score ranges to deployment tier recommendations.

Table 2: AEGF Deployment Recommendation by Score

AEGF Total Score	Deployment Recommendation
21–25	Tier 2–3 autonomous deployment; governance review at 90 days
16–20	Tier 1 co-pilot deployment; Tier 2 re-assessment at 6 months
11–15	Tier 0 supervised deployment only; address identified gaps first
≤ 10 or $G = 0$	Deployment not recommended; resolve fundamental readiness gaps

3.9 AEGF Architecture Overview

Figure 1 presents the integrated architecture of the AEGF, illustrating a three-layer governance structure: (i) the **Theoretical Foundation Layer** comprising NIST AI RMF, Sociotechnical Systems Theory, and the TOE Framework; (ii) the **Five-Dimension Assessment Layer** encompassing the scored dimensions and their sub-criteria; and (iii) the **Deployment Output Layer** articulating gate evaluation, tier assignment, and the progression review cycle. The interdependency gates (G_1 : Autonomy-Auditability; G_2 : Suitability-Readiness) are depicted as control-flow

connectors between Dimensions 2–3 and Dimensions 1–4 respectively, enforcing the constraint that composite score alone is insufficient for tier advancement.

Figure 1: AEGF Three-Layer Governance Architecture.




Figure 1. AEGF Three-Layer Architecture Diagram

4. Illustrative Application: AP Automation on Oracle ERP Cloud

4.1 Scenario and Deployment Context

Accounts payable processing on Oracle ERP Cloud represents one of the highest-priority agentic AI deployment candidates in the current generation of enterprise automation. The process is structurally rich — invoice ingestion, three-way match validation, approval routing, exception handling, and payment run initiation — and its exception component, historically the primary barrier to full RPA automation, is precisely the capability domain where agentic reasoning adds distinctive value. Oracle's Fusion Payables Agent, part of the March 2026 Fusion Agentic Applications suite [6], can retrieve contextual vendor history, assess dispute validity against contract terms, and route resolution recommendations through the approval workflow without a fixed decision tree governing each step.

4.2 AEGF Assessment

Dimension 1 — Process Suitability: The standard three-way matching process receives a 4/5 score for process structure, per-transaction financial impact, and regulation sensitivity being limited to the standard AP compliance rules. The process is scored a 3/5 on exception handling because of a low level of predictability of the structure. Suitability weighting assessment: HIGH for standard sub-process; MODERATE for exception sub-process.

Dimension 2 — Autonomy Tiering: According to Dimension 1, the standard invoice processing will be given Tier 1 (co-pilot) for the first implementation phase while the Tier 2 (captain) status will be reviewed in 90 days after the start date if anomaly levels stay reasonable. Exception handling starts at Tier 0.

Dimension 3 — Governance and Auditability: The natural audit framework provided by Oracle ERP Cloud along with OIC execution logging and Oracle Cloud Observability offers the sub-task-level tracing ability required by this dimension. Rollback options are provided by Oracle's Payment Cancellation and Journal Reversal processes. Score: 4/5.

Dimension 4 — Organisational Readiness: This dimension reflects the most significant governance gap at the current stage of Oracle Fusion agentic adoption. Formal AI governance roles are absent from most enterprise IT structures, AI literacy among ERP process owners is limited, and the Fusion Agentic Applications platform is sufficiently recent that even experienced Oracle implementation partners are developing their deployment methodologies concurrently with client engagements [20]. Score: 3/5 — targeted capability development required before advancing beyond Tier 1.

Dimension 5 — Risk and Continuity: Execution risk is mitigated through Oracle ERP Cloud's payment threshold and workflow configurations. Dependency risk is mitigated by using OIC's fault handling process. Auditability risk is minimal given the infrastructure considered in Dimension 3. Score: 4/5.

4.3 AEGF Score and Recommendation

Dimension	Score (1–5)
1. Process Suitability	4
2. Autonomy Tiering	4
3. Governance & Auditability	4
4. Organisational Readiness	3
5. Risk & Continuity	4
AEGF Total	19/25

Gate evaluation: $G(D_2, D_3) = G(4, 4) = 1 \checkmark$ (Autonomy-Auditability gate passes). $G(D_1, D_4) = G(4, 3) = 1 \checkmark$ (Suitability-Readiness gate passes; D_4 meets the minimum threshold of 3.0). Effective score $S^* = 19 \times 1 \times 1 = 19$.

A score of 19 maps to a Tier 1 Co-Pilot Deployment recommendation. Approval of the automation agent to perform autonomously in terms of the standard invoice processing in controlled financial parameters is given. The suggestions for exception handling should be reviewed by humans. The evaluation of AEGF in regard to Dimension 4 after six months of its launch is essential since its accomplishment in Tier 2 is contingent upon AI literacy and governance. Tier 2 status will be reviewed at 90 days provided anomaly rates remain within the $\epsilon \leq 0.02$ threshold.

4.4 Governed vs. Non-Governed Deployment

Table 3: AEGF-Governed vs. Non-Governed AP Agent Deployment

Governance Aspect	AEGF-Governed	Non-Governed
Autonomy boundary	Tier 1; defined financial thresholds	Undefined; full agent capability active
Audit trail	Sub-task level; OIC + Oracle Audit	Transaction level only; reasoning absent
Failure response	Graceful degradation to manual process	Ad hoc; recovery path undetermined
Compliance posture	Documented governance decisions	Unauditable; regulatory exposure elevated
Advancement pathway	Evidence-based, scheduled re-assessment	No formal review cycle

5. Discussion

5.1 Theoretical Contributions

Three distinct theoretical contributions emerge from this work. The first concerns the role of Sociotechnical Systems Theory in enterprise AI governance. Prior governance frameworks for autonomous AI — the NIST AI RMF [2], Acharya's maturity model [16], Otunlape's regulated-environment framework [7] — treat governance as primarily a technical and procedural design problem. By grounding AEGF Dimension 4 in Bostrom and Heinen's joint optimisation principle [4], the framework makes explicit what the technical-only framing obscures: governance failure in agentic ERP deployments is as likely to originate in institutional capacity gaps, misaligned accountability structures, and AI literacy deficits as in technical misconfiguration. Amanollahnejad et al.'s [21] empirical finding that sociotechnical alignment is the primary determinant of enterprise AI adoption outcomes provides direct support for this theoretical positioning.

The second contribution involves the repositioning of the TOE framework from an adoption prediction instrument to an active governance design variable. The existing cloud ERP adoption literature [17], [18] applies TOE to explain technology adoption decisions. The AEGF embeds the TOE's organisational and environmental dimensions within a scored deployment assessment instrument, treating them as first-class governance criteria rather than contextual background. This extension broadens the TOE's applicability to the ongoing governance of live AI deployments, not merely their initial adoption.

The third contribution is the provision of an operationalised, scored governance instrument that produces specific deployment tier recommendations. Saini [22] and Acharya [16] both articulate the need for practical agentic AI governance structures in enterprise contexts; neither provides an instrument with sufficient specificity for direct application to individual deployment decisions. The AEGF addresses this operationalisation gap.

5.2 Practical Implications

Enterprise architects working within Oracle ERP Cloud implementations can embed the AEGF assessment within the standard project governance cadence — during the solution design phase to establish the initial tier assignment, and at the transition to production to confirm that all five dimensions meet the minimum score threshold for the assigned tier. McKinsey's 2025 State of AI report finds that only 1% of organisations report achieving the full value anticipated from AI deployments at scale [8], underscoring that governance capacity — not just technical deployment capability — is the binding constraint on realising AI value at enterprise scale. Studies suggest that up to 70% of digital transformation initiatives fail to meet their stated objectives [8], providing a risk-based justification for the governance investment the AEGF requires.

For internal audit and compliance functions, the Governance and Auditability dimension (Dimension 3) provides a concrete specification of the audit trail, anomaly detection, and rollback requirements that should be verified before a deployment proceeds. The comparison in Table 3 makes the governance exposure differential between governed and ungoverned approaches concrete enough to support risk-based investment decisions at the senior leadership level.

5.3 Limitations

The AEGF has been developed as a conceptual framework and illustrated through a single representative deployment scenario; empirical validation across multiple organisations and platforms has not been conducted. The scoring scale reflects expert judgment rather than statistically validated criteria weights. The Oracle ERP Cloud capability references are current as of March 2026 and will require updating as the Fusion Agentic Applications platform continues to evolve. The framework's primary design reference is Oracle ERP Cloud; its applicability to SAP S/4HANA Cloud and Microsoft Dynamics 365 requires separate assessment.

5.4 Future Research Directions

Empirical validation of the AEGF across multiple organisational contexts and ERP platforms would provide the statistical evidence base needed to assess the relative predictive weight of each dimension. Longitudinal case studies tracking governance outcomes from initial Tier 0 deployment through Tier 2 or Tier 3 advancement would generate evidence about the conditions under which autonomy tier progression is achievable and sustainable. Extension of the framework to SAP and Microsoft Dynamics environments, and the development of a real-time dynamic AEGF score that incorporates ongoing execution monitoring data, represent technically tractable and practically consequential future directions.

6. Conclusion

Enterprise cloud ERP systems are deploying autonomous AI agents at a pace that governance practice has not kept up with. Gartner projects that 40% of enterprise applications will feature task-specific AI agents by 2026, up from less than 5% in 2025 [5]. Oracle's production launch of more than 600 Fusion agentic applications [6] means that the governance gap is no longer a future concern — it is a present one, and the consequences of ungoverned agentic deployment in live ERP environments are materially significant.

The Agentic ERP Governance Framework proposed here offers a structured, theoretically grounded, and practically executable response to that gap. Five dimensions — Process Suitability, Autonomy Tiering, Governance and Auditability, Organisational Readiness, and Risk and Continuity Management — form an integrated governance architecture that produces specific, defensible, evidence-based deployment recommendations. The accounts payable application on Oracle ERP Cloud demonstrates that the framework operates within realistic deployment constraints, producing governance outputs that are actionable without requiring specialised AI governance expertise beyond the reach of a senior enterprise architect.

Governance frameworks do not create value on their own. Their value is the structured assurance they provide that the capabilities they govern are deployed in ways that organisations can actually oversee, correct, and learn from. As agentic AI capabilities in cloud ERP continue to advance, that assurance will be one of the more consequential contributions a governance researcher can offer.

Declaration: This work represents the original contribution of the author and has not been submitted elsewhere for publication.

Funding: The author received no specific funding for this research.

Acknowledgement: The author acknowledges the support of Sri Venkateswara University, Tirupati, India.

Conflict of Interest: The author declares no conflict of interest.

Ethical Approval: This research did not involve human subjects, animal studies, or sensitive data.

Data Availability Statement: No new datasets were generated. All referenced data are available from the cited published sources.

REFERENCES

1. S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, "ReAct: Synergizing Reasoning and Acting in Language Models," in Proc. Int. Conf. Learning Representations (ICLR), 2023, arXiv:2210.03629.
2. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), U.S. Dept. of Commerce, Jan. 2023.
3. L. G. Tornatzky and M. Fleischer, *The Processes of Technological Innovation*, Lexington Books, 1990.
4. R. P. Bostrom and J. S. Heinen, "MIS Problems and Failures: A Socio-Technical Perspective. Part I: The Causes," *MIS Quarterly*, vol. 1, no. 3, pp. 17–32, 1977.
5. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," Gartner Press Release, Aug. 2025.
6. Oracle Corporation, "Oracle Fusion Agentic Applications," Oracle Product Announcement, Mar. 2026.
7. O. Otunla, "Governance Frameworks for Enterprise AI Systems Operating in Regulated Environments," *Int. J. Computer Applications*, vol. 187, no. 74, 2026, doi: 10.5120/ijca2026926244.

8. McKinsey & Company, The State of AI: How Organizations Are Rewiring to Capture Value, McKinsey Global Institute, 2025.
9. Bain & Company, "Is Agentic AI the Inflection Point for Scaling ERP Transformations?" Bain & Company Insights, 2025.
10. S. Sarferaz, "Implementing Agentic AI into ERP Software," ResearchGate Preprint, 2025.
11. J. Siderska et al., "Towards Intelligent Automation (IA): Literature Review on the Evolution of Robotic Process Automation (RPA), Its Challenges, and Future Trends," *Engineering Management Journal*, vol. 35, no. 4, 2023, doi: 10.2478/emj-2023-0030.
12. W. M. P. van der Aalst, M. Bichler, and A. Heinzl, "Robotic Process Automation," *Business & Information Systems Engineering*, vol. 60, no. 4, pp. 269–272, 2018, doi: 10.1007/s12599-018-0542-4.
13. M. M. Li et al., "Agentic AI," *Business & Information Systems Engineering*, 2026, doi: 10.1007/s12599-026-01009-w.
14. Y. K. Dwivedi et al., "Artificial Intelligence (AI): Multidisciplinary Perspectives on Emerging Challenges, Opportunities, and Agenda for Research, Practice and Policy," *Int. J. Information Management*, vol. 57, art. 101994, 2021, doi: 10.1016/j.ijinfomgt.2019.08.002.
15. Cloud Security Alliance, "Governance and NIST AI Agent Standards for Agentic Governance," CSA Research Paper, Mar. 2026.
16. V. Acharya, "Governing the Agentic Enterprise: A Governance Maturity Model for Managing AI Agent Sprawl in Business Operations," arXiv:2604.16338, Mar. 2026.
17. H. Hahoang, "Navigating the Digital Landscape: An Exploration of the Relationship Between Technology-Organization-Environment Factors and Digital Transformation Adoption in SMEs," *SAGE Open*, 2024, doi: 10.1177/21582440241276198.
18. U. Damali, M. Kocakulah, and A. S. Ozkul, "Investigation of Cloud ERP Adoption in the Healthcare Industry Through TOE Framework," *IGI Global*, 2023.
19. Y. Al-Slais and M. Ali, "Robotic Process Automation and Intelligent Automation Security Challenges: A Review," *IEEE/ResearchGate*, 2023.
20. K. Kirkpatrick, "Oracle's Fusion Agentic Apps: Can Platform-First AI Finally Deliver Enterprise ROI?" *Futurum Research*, 2026.
21. A. Amanollahnejad et al., "Aligning Socio-Technical Systems: Rethinking AI Adoption and Digital Transformation in SMEs," *Information Systems Management*, 2025, doi: 10.1080/10580530.2025.2612175.
22. S. Saini, "Governing the Agentic Enterprise: A New Operating Model for Autonomous AI at Scale," *California Management Review Insights*, Mar. 2026.