

A Risk-Tiering Framework for Governing SAP Generative and Agentic Artificial Intelligence in Regulated Enterprise Environments: A Utility Sector Case Study

Pon Murugesh Devendren

Independent Researcher, USA. ORCID: 0009-0007-9554-8229

Abstract: Generative and agentic artificial intelligence systems are entering production on SAP Business Technology Platform at regulated enterprises whose governance frameworks were not designed for them. This paper presents a risk-tiering framework for governing SAP generative and agentic AI deployments in regulated enterprise environments, developed through design science methodology and validated through a case study in the utility sector. The framework defines four risk tiers: minimal-risk assistive AI, moderate-risk recommendatory AI, high-risk autonomous AI with consequential business impact, and critical-risk AI in safety or regulatory reporting, mapped to SAP process domains and data sensitivity classifications. Governance controls at each tier are calibrated to the autonomy level and regulatory exposure of the AI use case. The framework is aligned by construction to the risk classification structure of the EU AI Act, the NIST AI Risk Management Framework, and SAP's ISO/IEC 42001-certified responsible AI posture. A case study in the utility sector involving 15 generative and agentic AI use cases across five SAP domains demonstrates practical application, producing a tier distribution of approximately 30% Tier 1, 40% Tier 2, 25% Tier 3, and 5% Tier 4 and a governance operating model reviewed by external auditors and regulatory affairs professionals. The framework provides enterprise architects and AI governance leads in regulated SAP environments with an implementable, auditor-defensible governance instrument that is calibrated to the SAP platform context.

Keywords: AI governance, Risk Tiering, SAP, Agentic AI, EU AI Act, NIST AI RMF, Regulated Industries, Enterprise AI

1. INTRODUCTION

Regulated enterprises have been running SAP as their system of record for decades. What is new is the embedding of generative and agentic AI directly into the SAP process landscape, which existing governance frameworks were not built to handle. Utilities, financial institutions, and healthcare organizations now face the prospect of AI agents approving invoices, updating asset records, preparing journal entries, and drafting regulatory submissions inside the same SAP environment that feeds their external reporting, their audit evidence, and their safety-critical operations [26]. Generic AI policy frameworks [11], [13] operate at a level of abstraction that does not translate into the specific process decisions a SAP enterprise architect must make in 2026. The EU AI Act [6], reaching full enforcement for high-risk AI systems on August 2, 2026, provides legally binding risk categories that were not designed for the SAP process taxonomy. The NIST AI RMF [3] provides a comprehensive risk management vocabulary without mapping its functions to specific SAP deployment decisions.

The academic literature confirms the gap rather than filling it. Research on AI governance broadly [18], [19] establishes that governance is necessary and characterizes its dimensions but does not produce the SAP-specific instrument enterprises need. Research on AI in enterprise information systems [24], [25] and AI in the energy sector [27] provides relevant context without resolving the classification question: which use case requires which controls, and on what basis? AI governance research for the public sector [16], [28] offers useful parallels in a different



institutional context. The governance challenge is sharpened by the emergence of agentic AI architectures, systems perceiving their environment, reasoning over goals, and executing autonomous action sequences [15], whose accountability structures existing IT governance frameworks [9] were not designed to accommodate.

1.1 Research Contributions

This paper addresses the governance gap through four contributions to the literature on enterprise AI governance and information systems:

Contribution 1: SAP-calibrated four-tier risk classification framework. Unlike the NIST AI RMF [3], which provides risk management functions without process-specific use case classification, the proposed framework applies sequential classification criteria, SAP write access, data sensitivity, and regulatory reporting connection, derived from the SAP process taxonomy. This produces tier assignments that are directly actionable for SAP enterprise architects rather than requiring further translation from a technology-agnostic framework.

Contribution 2: Tier-specific governance controls traceable to regulatory articles. In contrast to generic enterprise AI governance approaches [18], [19], which specify governance dimensions but do not provide regulatory article-level traceability, every control in the proposed framework is explicitly traceable to a specific EU AI Act article [6], NIST AI RMF function [3], or SAP certification requirement [1]. This traceability is what makes the framework auditor-defensible rather than merely internally coherent.

Contribution 3: SAP-specific governance operating model. Unlike ISO/IEC 42001 [1], the standard-clause level without use-case-level implementation guidance, the proposed governance operating model specifies how the model specifies how SAP AI Agent Hub [5] serves as the central visibility and monitoring layer, how the use case intake process should be designed for business-unit accessibility, and how the governance review board should be structured for cross-functional decision authority.

Contribution 4: Utility sector case study with primary deployment metrics. Unlike existing regulated industry AI governance literature, which is primarily conceptual [27], [28], the utility sector case study provides primary deployment evidence from a four-month engagement involving 15 use cases across five SAP domains, seven cross-functional stakeholder groups, and an approval timeline differential ranging from four to six weeks for Tier 1 and Tier 2 to approximately four months for Tier 4.

2. METHOD

2.1 Design Science Approach

The framework was developed using design science methodology [8], the appropriate choice for IS research producing a prescriptive artifact intended to solve an identified organizational problem. Design science methodology requires explicit treatment of three spaces: the problem space, the design space, and the evaluation space. The problem space was characterized through analysis of three converging constraints: the SAP process landscape in regulated industries, the data sensitivity stratification requirements of those industries, and the regulatory obligations imposed by the EU AI Act [6], NIST AI RMF [3], and sector-specific frameworks. The design space drew on AI ethics frameworks [11], [20], EU regulatory guidance [6], [10], AI risk taxonomies [14], [23], IS governance theory [9], and AI governance research [18], [19], [29]. The evaluation space was the utility sector case study, a real organizational portfolio of 15 AI use cases subjected to the framework with cross-functional stakeholder review over a four-month period, including external auditors and regulatory affairs professionals.

Stakeholder engagement was structured across three phases. In the problem characterization phase, semi-structured interviews were conducted with representatives from Information Technology, Information Security, Legal and Compliance, and Internal Audit to map existing governance gaps against the emerging SAP AI deployment pipeline. In the design validation phase, two facilitated workshops were held with the full seven-group cross-functional stakeholder set to review draft tier definitions and classification criteria, with business unit sponsors providing domain-

specific use case context and compliance stakeholders stress-testing regulatory traceability claims. In the evaluation phase, external auditors reviewed the Tier 4 control specifications and audit evidence framework against their existing internal control testing methodology, and regulatory affairs professionals assessed whether the Tier 3 and Tier 4 control sets adequately reflected FERC and EPA reporting obligations. Classification decisions were validated iteratively: each use case owner reviewed the tier assignment rationale and had the opportunity to contest classification before the governance review board issued a final determination.

2.2 Framework Design Principles

Calibration to SAP process criticality holds that tier classification must reflect the consequences of AI failure within SAP's process architecture, not generic AI risk categories. A procurement agent in SAP Ariba carries different downstream risk than a document summarization agent operating against SAP Help Portal content, even when both use the same underlying model. The difference lies in the process consequences of AI error, not in the AI technology employed.

Data sensitivity primacy holds that the sensitivity classification of the SAP data the AI system accesses is a more reliable tier determinant than the sophistication of the AI technology. A rule-based automation processing financial reporting data requires more governance rigor than a sophisticated large language model answering questions about publicly available product documentation. This inverts the common organizational intuition that more sophisticated AI requires more governance, a tendency the algorithmic accountability literature identifies as a systematic source of governance failure [17].

Regulatory traceability holds that every tier boundary and governance control must be traceable to a specific regulatory article, standard function, or certification requirement. Internal governance coherence is insufficient if the framework cannot demonstrate alignment to the external obligations regulators and auditors will test against [2], [6]. Traceability is the property that distinguishes an auditor-defensible governance framework from one designed for internal consumption.

Minimum sufficiency holds that controls at each tier must be the minimum necessary to address the risk profile of that tier. The principle applies in both directions: under-controlling high-risk use cases creates regulatory exposure, while over-controlling low-risk ones creates compliance fatigue that discourages responsible AI adoption [4], [7].

2.3 Regulatory Alignment Methodology

The framework's regulatory alignment operates at three layers. The EU AI Act layer maps Tier 1 and Tier 2 to the Act's minimal and limited risk categories; Tier 3 and Tier 4 correspond to the Annex III high-risk classification covering critical infrastructure, essential services, and systems influencing decisions with significant consequences for individuals [6]. The NIST AI RMF layer maps each governance operating model component to an RMF function [3]: Govern maps to the review board and decommissioning procedures. The map corresponds to the intake and tier classification processes. Measure maps to ongoing monitoring and performance metrics; Manage maps to circuit breaker controls, the Tier 3 supervisor cadence, and the Tier 4 parallel human process requirement. The SAP responsible AI layer positions the customer-side framework as the use-case-level complement to SAP's ISO/IEC 42001 certification [1] and NIST CSF Tier 3 alignment [22], building on the platform baseline rather than duplicating it.

Table I. Comparison of Proposed Framework Against Existing Governance Approaches Across Seven Dimensions

Dimension	Proposed Framework	NIST AI RMF [3]	EU AI Act [6]	ISO/IEC 42001 [1]	Generic Enterprise AI Gov.
SAP process specificity	Full (calibrated to SAP)	None (technology-agnostic risk functions)	None (general-purpose risk categories)	None (platform-agnostic)	None (vendor-agnostic checklists)

	process taxonomy and data sensitivity classifications)			standard clauses)	
Tier granularity	Four tiers with sequential, criterion-based decision gates	Risk levels undefined (Govern/Map/Measure/Manage functions only)	Minimal / Limited / High / Unacceptable (broad categories)	General risk categories without tier boundaries	Typically binary: high risk / low risk
Regulatory traceability	Every control traceable to specific EU AI Act article, NIST function, or SAP certification clause	Functions specified; articles not mapped to controls	Articles specified; operationalization left to implementer	Standard clauses general; no article-level mapping	Rarely mapped to specific regulatory articles
Agentic AI coverage	Explicit (AI autonomy level is a primary tier classification criterion)	Emerging guidance only; agentic AI not addressed specifically	High-risk category covers some autonomous use cases	Limited coverage of agentic AI deployment specifics	Rarely addresses agentic AI autonomy explicitly
Minimum sufficiency principle	Explicit design rule: controls at each tier are the minimum necessary to address that tier's risk profile	Not specified (full Govern/Map/Measure/Manage applies uniformly)	Not specified (high-risk obligations apply uniformly)	Not specified (standard clauses apply uniformly)	Not specified (controls typically uniform across risk levels)
Auditability provisions	Tier-specific audit evidence requirements; Tier 4 requires full output lineage and independent assessor	Measure and Manage functions address monitoring; evidence not specified	Article 17 quality management system; evidence requirements general	Clause 9 performance evaluation; audit requirements at clause level	Variable (audit provisions depend on specific framework chosen)

Implementati on guidance	SAP- specific operating model: intake form, review board, SAP AI Agent Hub monitoring [5]	Process guidance; no platform-specific implementation	Obligation statements; compliance pathway left to implementer	Standard clauses; implementati on guidance through conformity assessment	Generic checklists; no platform- specific operationalizati on
-----------------------------	---	---	---	--	--

The seven-dimension comparison in Table I confirms that the existing landscape offers no single instrument combining SAP process specificity, four-tier granularity, article-level regulatory traceability, explicit agentic AI coverage, and minimum sufficiency design. Each existing approach addresses some dimensions but not all simultaneously. The proposed framework resolves all seven dimensions for the SAP regulated-industry context, the combination that the practical governance challenge requires. Table II summarizes how each design principle maps to specific framework features and the sections in which it is operationalized.

Table II. Mapping of Framework Design Principles to Framework Features and Operationalization

Design Principle	Framework Features Driven	Where Operationalized
Calibration to SAP process criticality	Four-tier structure; sequential classification criteria (write access, data sensitivity, regulatory connection); tier boundary definitions	Section 3.1; Table III (tier classification); utility sector use case examples per tier
Data sensitivity privacy	Data sensitivity as second sequential classification gate; sensitivity tier labels (Public, Internal/Confidential, Restricted/Regulated) embedded in tier definitions	Section 3.1; Table III (data sensitivity column); intake form design in Section 4.1
Regulatory traceability	Per-control EU AI Act article citations (Articles 9, 13, 14, 17); NIST AI RMF function mapping per operating model component; SAP ISO/IEC 42001 clause alignment	Section 2.3; Table I (regulatory comparison); Section 4.2
Minimum sufficiency	Prohibition on applying higher-tier controls to lower-tier use cases; lightweight Tier 1 controls; progressive control addition across tiers	Sections 3.1.1 through 3.1.4; Table IV (governance controls by tier)

3. RESULTS

3.1 The Four-Tier Risk Classification Framework

The framework classifies SAP generative and agentic AI use cases into four tiers using three primary criteria applied sequentially as decision gates: SAP system write access, data sensitivity classification, and regulatory reporting connection. A use case is evaluated against each criterion in order; escalation is triggered by the first criterion that applies. Table III presents the four-tier structure.

Table III. Four-Tier Risk Classification Framework for SAP AI Deployments

Tier	Label	Autonomy Level	SAP Write Access	Data Sensitivity	Regulatory Connection	Governance Intensity
1	Minimal Risk	Assistive / Informational	None	Public / General	None	Lightweight
2	Moderate Risk	Recommendatory	None (human triggers)	Internal / Confidential	Indirect	Structured review
3	High Risk	Autonomous (bounded)	Yes, within limits	Confidential / Restricted	Upstream feed	Substantial
4	Critical Risk	Autonomous / semi-autonomous	Yes	Restricted / Regulated	Direct	Maximum

3.1.1 Tier 1: Minimal Risk: Assistive and Informational AI

Tier 1 covers AI use cases where the system provides information or suggestions to a human user who retains complete decision-making authority, the underlying SAP data is non-sensitive, and the downstream consequences of AI error are limited to inconvenience or rework. Three sequential criteria must all evaluate negatively for Tier 1 classification: no write access to any SAP system of record; no data accessed above the public or general-use sensitivity tier; and no regulatory reporting obligation connected to any data or process the AI touches. Utility sector examples include Joule-powered natural language queries against publicly reportable financial summaries, AI-assisted generation of training documentation from SAP Help Portal content, and conversational search over SAP knowledge articles for IT support queries.

Governance at Tier 1 is lightweight: use case registration in SAP AI Agent Hub [5], standard role-based access control applied to the underlying data, and basic audit logging. Human oversight is passive; users can review and disregard AI outputs without consequence. The minimum sufficiency principle prohibits imposing Tier 3 controls on Tier 1 use cases; the compliance burden would be disproportionate and would undermine adoption of AI use cases carrying negligible governance risk [4].

3.1.2 Tier 2: Moderate Risk: Recommendatory AI with Human Decision

Tier 2 covers AI use cases where the system generates a recommendation, ranking, or draft output that a human reviews and accepts, modifies, or rejects, with the human decision triggering any SAP write action. Utility examples include AI-generated invoice matching recommendations in SAP Ariba reviewed by a procurement analyst, Joule-assisted draft responses to regulatory information requests reviewed by a compliance manager, and AI-ranked vendor shortlists reviewed by a category manager before initiating sourcing events.

The Tier 2 governance principle is that human review must be substantive rather than procedural. A rubber-stamp approval process defeats the governance purpose; the reviewer bears accountability for the decision without having exercised the judgment that accountability implies, creating liability without the oversight that liability is meant to accompany. Designing substantive review requires presenting the reviewer with the AI's reasoning alongside its recommendation: the key data points, assumptions, and confidence indicators that produced it so the reviewer can evaluate whether the recommendation is appropriate for the specific context [21]. This explainability requirement operationalizes EU AI Act Article 13's transparency obligations [6] and directly addresses the accountability gap the algorithmic decision-making literature identifies as a systemic governance failure mode [17].

3.1.3 Tier 3: High Risk: Autonomous AI with Consequential Business Impact

Tier 3 covers AI use cases where the system executes actions in SAP systems of record autonomously, generating financial transactions, updating master data, triggering procurement workflows, or modifying operational

configurations within defined parameters without requiring human confirmation for each individual action. Utility examples include autonomous invoice processing agents matching, approving, and posting invoices to SAP Financial Accounting within defined financial limits; agentic procurement workflows creating and submitting purchase orders up to a defined threshold; and AI systems updating SAP Asset Management records from field sensor data without manual review.

Governance controls at Tier 3 are substantially more demanding than Tier 2: a formal risk assessment documented against the NIST AI RMF [3] and EU AI Act Article 9 [6]; a defined SAP authorization scope specifying the exact transaction boundary of the AI agent; a mandatory human supervisor review cadence for autonomous AI actions; circuit breaker controls automatically suspending agent execution when anomaly thresholds are breached; and a rollback procedure documented and tested before production deployment. A sub-classification within Tier 3 distinguishes well-bounded autonomous processing with no downstream regulatory reporting connection from autonomous processing in domains that feed upstream into regulatory reporting chains, the latter requiring additional pre-reporting human review and a documented process owner attestation.

3.1.4 Tier 4: Critical Risk: AI in Safety, Compliance, or Regulatory Reporting

Tier 4 covers AI use cases that directly generate, validate, or influence outputs used in external regulatory submissions, safety-critical operational decisions, or financial statements subject to audit. In the utility sector, this tier includes AI systems processing data feeding into FERC reporting, EPA emissions reporting, and SAP Financial Accounting financial statement preparation and AI-assisted analysis supporting safety-critical grid management or asset maintenance decisions.

The control set at Tier 4 is the most stringent: independent technical review by a qualified assessor outside the development team; documented alignment to all applicable regulatory requirements; a parallel human process that can produce the same output independently for validation; mandatory explainability tracing every AI-generated output to the specific SAP data and reasoning steps that produced it [21]; and a formal decommissioning procedure. EU AI Act Articles 9, 13, 14, and 17, covering risk management systems, transparency, human oversight, and quality management, are collectively operationalized through this control set [6].

3.2 *Utility Sector Case Study*

3.2.1 Deployment Context

The case study involved designing and implementing a governance framework for AI at a large North American utility operating SAP S/4HANA across finance, procurement, asset management, and workforce operations. The organization had identified 15 generative and agentic AI use cases across five SAP domains, Finance, Procurement, Asset Management, Human Capital Management, and Integration Suite, with no existing AI governance process. The regulatory environment included FERC oversight, state public utility commission requirements, and EPA emissions reporting obligations. The engagement involved seven cross-functional stakeholder groups: Information Technology, Information Security, Legal, Compliance, Internal Audit, Finance, and business unit sponsors, over a four-month period. Stakeholder involvement followed a structured engagement model. Business unit sponsors participated in intake workshops where use case submissions were developed collaboratively using the intake form described in Section 4.1, ensuring that classification inputs reflected operational process knowledge rather than sponsor assumptions about AI risk. Cross-functional stakeholder groups reviewed tier assignments and proposed control specifications in governance board sessions held on a two-week cadence for Tier 1 and Tier 2 use cases and monthly for Tier 3 and Tier 4. External auditors engaged at two points: a mid-engagement review of draft Tier 4 control specifications to assess auditability, and a final review of the completed audit evidence framework against the organization's existing SOX internal control documentation. Regulatory affairs professionals provided written assessments of whether AI use in FERC and EPA reporting processes required proactive regulatory disclosure, and these assessments directly informed the Tier 4 regulatory disclosure control in Table IV. All client-identifying information has been generalized in accordance with confidentiality obligations.

3.2.2 Use Case Portfolio Assessment and Tier Distribution

Applying the sequential classification criteria of Section 3.1 to the full portfolio of 15 use cases produced a tier distribution of approximately 30% Tier 1, 40% Tier 2, 25% Tier 3, and 5% Tier 4. The high proportion of Tier 2 use cases reflected the utility sector's natural risk conservatism: many business sponsors who initially requested fully autonomous AI processing accepted a recommendatory model once they understood that Tier 2 controls could be designed for fast, low-friction human review rather than bureaucratic approval processes. Tier 3 use cases concentrated in procurement and accounts payable, where the financial controls framework was mature and the transaction boundary well-defined enough to make autonomous processing within limits defensible. Only two use cases, AI-assisted journal entry preparation and account reconciliation in SAP Financial Accounting, were classified as Tier 4, representing the 5% share of the portfolio where the direct regulatory reporting connection triggered the most stringent controls.

The Tier 4 approval process required approximately four months compared to four to six weeks for Tier 1 and Tier 2, consistent with the framework's intent that Tier 4 approval should be deliberately rigorous. The extended duration reflected the depth of work required: external auditor engagement to define acceptable evidence of AI oversight, regulatory affairs coordination to assess whether AI use in financial reporting required regulatory disclosure, and formal legal opinion on whether the existing internal control framework adequately covered AI-generated journal entries.

3.2.3 Governance Framework Artifacts

The engagement produced five governance artifacts: (i) the use case intake form, designed collaboratively with business sponsors in business process language without requiring specialist AI or governance knowledge; (ii) the risk classification methodology applying the sequential criteria of Section 3.1; (iii) the governance review board charter defining membership, cadence, and decision authority; (iv) tier-specific control specifications including circuit breaker parameters for Tier 3 and the audit evidence framework for Tier 4; and (v) the ongoing monitoring framework integrating SAP AI Agent Hub [5] as the central visibility layer for deployed AI agent status.

Table IV. Governance Controls by Tier and Case Study Outcome Metrics

Control Dimension	Tier 1	Tier 2	Tier 3	Tier 4
Use case registry	(SAP AI Agent Hub)	Yes	Yes	Yes
Audit logging	Basic	Full	Full + anomaly detection	Full + tamper-evident
Human oversight	Passive	Structured review	Supervisor cadence	Parallel human process
Risk assessment	None	Lightweight	NIST AI RMF documented	NIST AI RMF + EU AI Act Art. 9
Circuit breaker	None	None	Required	Required
Explainability	None	Recommendation context + key data points	Full action trace	Full output lineage
External review	None	None	None	Independent assessor required
Regulatory disclosure	None	None	None	Formal assessment required

Rollback procedure	None	None	Documented and tested	Documented and tested
Approval timeline (case study)	4–6 weeks	4–6 weeks	4–6 weeks	Approximately four months

The governance control progression across tiers in Table IV reflects the minimum sufficiency principle: each tier adds only the controls its risk profile requires. The approximately four-month Tier 4 approval duration, substantially longer than the four- to six-week timeline for Tier 1 through Tier 3, reflects the deliberate rigor the framework specifies for AI use cases directly connected to regulatory submissions and audited financial statements. The governance artifacts produced at Tier 4, specifically the audit evidence framework and the regulatory disclosure assessment, subsequently served as reference designs for other regulated utility organizations in the region, indicating that the governance investment at the most stringent tier produces durable, transferable outputs rather than one-time compliance documentation.

4. DISCUSSION

4.1 Governance Operating Model

The framework is operationalized through three components that translate the tier classification structure into a functioning organizational governance process. The use case intake process is a structured submission mechanism that collects the SAP process domain, specific SAP data access requirements (including whether write access is required), regulatory reporting chain connections, proposed human oversight mechanisms, and the business value case. Designing the intake form collaboratively with business sponsors, framing questions in business process terms with guidance notes explaining why each piece of information was required, produced returns in submission quality: the proportion of submissions requiring clarification requests dropped after the second governance board cycle, and classification decision quality improved because the information provided was more accurate and complete.

The governance review board brings cross-functional authority to each classification decision. Permanent members from IT Architecture, Information Security, Legal and Compliance, and Internal Audit provide consistent governance rigor; business unit representatives attend for use cases from their domain, ensuring that domain context informs each decision without requiring all board members to hold it. A two-tier cadence, two weeks for Tier 1 and Tier 2, monthly for Tier 3 and Tier 4, balances review thoroughness against business momentum, with an expedited track requiring written justification and deputy CIO approval for urgent business cases. The ongoing monitoring framework uses SAP AI Agent Hub [5] as the central visibility layer for all deployed AI agents, with automated anomaly alerts reducing the burden of manual log review and aligning directly to the NIST AI RMF's Measure and Manage functions [3].

4.2 Regulatory Alignment

The August 2, 2026 enforcement of the EU AI Act's high-risk obligations [6] converts previously voluntary governance aspirations into legally binding requirements with material financial penalties for non-compliance [2]. Regulated enterprises that have deferred AI governance as a strategic conversation are facing a compressing window of non-compliance that reactive approaches cannot close quickly. Tier 3 and Tier 4 controls operationalize Articles 9 (risk management systems), 13 (transparency), 14 (human oversight), and 17 (quality management) in a form immediately reviewable against the Act's text, an enterprise demonstrating this mapping to a regulator or auditor has a materially stronger governance posture than one citing general RMF alignment without use-case-level operationalization [29], [30].

The NIST AI RMF mapping [3] links each operating model component to a specific RMF function, enabling enterprises to demonstrate RMF alignment at the use-case level rather than as a general organizational commitment. SAP's ISO/IEC 42001 certification [1] and NIST CSF Tier 3 alignment [22] provide the platform-level governance baseline on which the customer framework builds. Platform certification covers SAP's AI infrastructure, Joule, SAP

AI Core, and SAP AI Launchpad; it cannot cover how a specific customer classifies a specific use case against that customer's specific regulatory obligations. The customer-side framework closes this gap, not by duplicating platform-level assurance but by extending it to the use-case-level governance decisions platform certification cannot reach [31].

4.3 Agentic AI and IS Governance Implications

The shift from AI as decision support to AI as an organizational actor [24] is what makes tier-based governance necessary. An AI system that provides a recommendation is qualitatively different from one that executes a transaction; the latter's accountability structures cannot be handled by the same controls as the former. The tier classification approach makes this distinction explicit and derives control requirements from it, producing a governance framework that scales with AI autonomy [15] rather than applying uniform controls regardless of what the AI system actually does. This directly addresses the accountability challenge identified in IS research on agentic AI as a new frontier: that existing governance frameworks lack the structural vocabulary to distinguish AI systems by their capacity for consequential autonomous action [24].

The explainability requirements embedded in Tier 2 through Tier 4 [21] address the accountability gap the algorithmic decision-making literature identifies as a systemic governance failure [17]: when AI systems influence or make decisions with significant consequences, the system must produce explanations enabling human accountability to be exercised meaningfully. The utility sector case study demonstrates this requirement is operationally achievable within SAP environments; the Tier 2 procurement review interface presented structured explanations, including key data points and alternatives considered; and external auditors treated this as substantive governance evidence rather than mere documentation compliance. Seghid et al.'s emerging threats taxonomy [14] identifies uncontrolled autonomous action chains as a primary enterprise AI risk; Tier 3 and Tier 4 circuit breakers and bounded transaction scopes directly address this specific risk class [32].

4.4 Generalizability and Limitations

The framework's core design principles, calibration to process criticality, data sensitivity primacy, regulatory traceability, and minimum sufficiency are generalizable to any enterprise ERP environment deploying AI in regulated process domains [25], [26]. Tier boundaries and control specifications require adjustment for non-SAP platforms, but the classification logic and governance operating model are transferable. Comparative AI governance research confirms that regulated contexts require framework adaptation rather than wholesale adoption of generic instruments [28], [29], consistent with the design science approach taken here [8].

Several limitations apply. The case study represents a single organizational context, a large North American utility, and the tier distribution and governance timeline findings may not generalize to utilities of different sizes, regulatory profiles, or SAP implementation maturities. The framework's Tier 3 and Tier 4 controls are calibrated to the EU AI Act and NIST AI RMF as they stood at mid-2026; both frameworks are subject to ongoing guidance and interpretation that may require control updates. The normative AI ethics foundations of the framework [7], [11], [12], [20] are operationalized in its design principles but are not fully mapped to individual control specifications; future work should make these normative connections more explicit. As SAP's agentic capabilities evolve, particularly multi-agent orchestration through SAP AI Agent Hub [5] and Joule Studio, Tier 3 use cases manageable today may require Tier 4 controls as agent autonomy and delegation depth increase [33], [34].

5. CONCLUSION

The regulatory landscape for AI in regulated industries has changed materially. The EU AI Act's August 2, 2026, enforcement deadline, the NIST AI RMF's growing adoption as a regulatory reference standard, and SAP's ISO/IEC 42001 certification collectively establish an environment in which AI governance is no longer a voluntary best practice; it is an enforceable expectation with material legal consequences for non-compliance. The framework presented here gives regulated SAP enterprises a governance instrument calibrated to that environment: four risk tiers grounded in SAP process criticality and data sensitivity, governance controls traceable to specific regulatory articles

and standard functions, and a governance operating model validated through production deployment across a 15-use-case utility sector portfolio. The case study of the utility sector demonstrates that the framework is implementable within a realistic project timeline, four months for a complete portfolio assessment covering five SAP domains and seven cross-functional stakeholder groups, when governance design engages business sponsors, compliance stakeholders, and external auditors simultaneously rather than sequentially. The tier distribution produced, approximately 30% Tier 1, 40% Tier 2, 25% Tier 3, and 5% Tier 4, reflects a pattern likely to be representative of other regulated utility sector SAP environments, where business sponsor risk conservatism and mature financial controls in procurement and accounts payable combine to skew toward Tier 2 and concentrate Tier 3 use cases in financially bound autonomous processing. The framework's defining principle, that tier classification is driven by what data the AI system touches, what decisions it influences, and what regulatory obligations those decisions connect to, rather than by the sophistication of the AI technology employed, provides a governance vocabulary that resolves the practical decision regulated enterprises face: not whether to govern AI, but which use case gets which controls and why. Three directions warrant future investigation: a multi-organization cross-sector study testing the generalizability of tier distribution findings across financial services, healthcare, and public utilities; systematic evaluation of framework integration with SAP's emerging multi-agent orchestration capabilities, including SAP AI Agent Hub and Joule Studio, as those capabilities reach general availability; and extension of Tier 4 controls to address the sector-specific evidentiary requirements of FERC, OCC, and FDA, whose AI governance expectations are developing faster than cross-sector frameworks can accommodate.

REFERENCES

1. Vanessa Barber, Hedayatollah Hosseini, and Peter Westphal, "How SAP is safeguarding its customers: Implementing the NIST Cybersecurity Framework achieving Tier 3 alignment," SAP SE and NIST, 2024. Available: <https://www.sap.com/documents/2024/08/e49580a9-d27e-0010-bca6-c68f7e60039b.html>
2. Cloud Security Alliance AI Safety Initiative, "EU AI Act High-Risk Deadline: Enterprise Readiness Gap," Cloud Security Alliance, 2026. Available: <https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-compliance-deadline-20/>
3. National Institute of Standards and Technology, "AI Risk Management Framework v1.0," NIST, 2023. Available: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
4. Fabio Correa Xavier, "Governing frontier general-purpose AI in the public sector: Adaptive risk management and policy capacity under uncertainty through 2030," arXiv preprint arXiv:2604.06215, 2026. Available: <https://arxiv.org/abs/2604.06215>
5. SAP SE, "SAP AI Agent Hub," 2026. Available: <https://www.sap.com/products/artificial-intelligence/ai-agent-hub.html>
6. European Parliament and Council, "Regulation (EU) 2024/1689 — Artificial Intelligence Act," Official Journal of the EU, 2024. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689
7. Arjun Panesar, "Ethics of intelligence," in *Machine Learning and AI for Healthcare*, Apress, 2019, pp. 207–254. Available: https://link.springer.com/chapter/10.1007/978-1-4842-3799-1_6
8. Roel Wieringa, *Design Science Methodology for Information Systems and Software Engineering*. Springer, 2014. Available: <https://link.springer.com/book/10.1007/978-3-662-43839-8>
9. Steven de Haes and Wim Van Grembergen, "Enterprise IT governance, business/IT alignment and value creation," in *Proc. HICSS*, 2009. Available: <https://download.e-bookshelf.de/download/0003/9278/36/L-G-0003927836-0013313330.pdf>
10. Claudio Novelli et al., "Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity," *Computer Law & Security Review*, vol. 55, p. 106066, 2024. Available: <https://doi.org/10.1016/j.clsr.2024.106066>
11. Luciano Floridi et al., "AI4People — An ethical framework for a good AI society," in *Ethics, Governance, and Policies in Artificial Intelligence*, 2018, pp. 19–39. Available: <https://pubmed.ncbi.nlm.nih.gov/30930541/>
12. Jamie Berryhill et al., "Hello, World: Artificial intelligence and its use in the public sector," *OECD Working Papers on Public Governance*, 2019. Available: <https://oecd-opsi.org/wp-content/uploads/2019/11/AI-Report-Print.pdf>
13. Nathalie A. Smuha, "The EU approach to ethics guidelines for trustworthy artificial intelligence," *Computer Law Review International*, vol. 20, no. 4, pp. 97–106, 2019. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3443537
14. Niyat Seghid et al., "Emerging threats in AI: A detailed review of misuses and risks across modern AI technologies," *Frontiers in Communications and Networks*, vol. 6, p. 1727425, 2025. Available: <https://doi.org/10.3389/frcmn.2025.1727425>
15. Lei Wang et al., "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024. Available: <https://doi.org/10.1007/s11704-024-40231-1>

16. Luca Tangi, Colin Van Noordt, and A. Paula Rodriguez Müller, "The challenges of AI implementation in the public sector: An in-depth case studies analysis," in *Proc. 24th Annual International Conference on Digital Government Research*, 2023, pp. 414–422. Available: <https://doi.org/10.1145/3598469.3598516>
17. Bruno Lepri et al., "Fair, transparent, and accountable algorithmic decision-making processes," *Philosophy & Technology*, vol. 31, no. 4, pp. 611–627, 2018. Available: <https://doi.org/10.1007/s13347-017-0279-x>
18. Emmanouil Papagiannidis, Patrick Mikalef, and Kieran Conboy, "Responsible artificial intelligence governance: A review and research framework," *Journal of Strategic Information Systems*, vol. 34, no. 2, p. 101885, 2025. Available: <https://doi.org/10.1016/j.jsis.2024.101885>
19. Mark Anthony Camilleri, "Artificial intelligence governance: Ethical considerations and implications for social responsibility," *Expert Systems*, vol. 41, no. 7, p. e13406, 2024. Available: <https://doi.org/10.1111/exsy.13406>
20. Luciano Floridi et al., "How to design AI for social good: Seven essential factors," in *Ethics, Governance, and Policies in Artificial Intelligence*, Springer, 2021, pp. 125–151. Available: https://doi.org/10.1007/978-3-030-81907-1_9
21. Giulia Vilone and Luca Longo, "Notions of explainability and evaluation approaches for explainable artificial intelligence," *Information Fusion*, vol. 76, pp. 89–106, 2021. Available: <https://doi.org/10.1016/j.inffus.2021.05.009>
22. National Institute of Standards and Technology, "NIST Cybersecurity Framework 2.0," NIST CSWP 29, 2024. Available: <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>
23. Georgios Agathangelos-Ioannou and Aggeliki Tsohou, "Methodologies for information security risk assessment: Shortcomings for analysing artificial intelligence systems," *International Journal of Information Security*, vol. 25, no. 3, p. 98, 2026. Available: <https://doi.org/10.1007/s10207-026-01271-9>
24. Naveen Kumar, Xiahua Wei, and Han Zhang, "Agentic artificial intelligence as a new frontier in information systems: Promise, peril, and research opportunities," *Information & Management*, p. 104317, 2026. Available: <https://doi.org/10.1016/S0378720626000170>
25. Veda C. Storey et al., "Generative artificial intelligence: Evolving technology, growing societal impact, and opportunities for information systems research," *Information Systems Frontiers*, 2025. Available: <https://doi.org/10.1007/s10796-025-10581-7>
26. Luminita Hurbean et al., "Unveiling AI-powered ERP systems: A systematic review of trends and future directions," *Information Systems and e-Business Management*, 2026. Available: <https://doi.org/10.1007/s10257-026-00728-7>
27. Roba Alsaigh, Rashid Mehmood, and Iyad Katib, "AI explainability and governance in smart energy systems: A review," *Frontiers in Energy Research*, vol. 11, p. 1071291, 2023. Available: <https://doi.org/10.3389/fenrg.2023.1071291>
28. Khalifa Alhosani and Saadat M. Alhashmi, "Opportunities, challenges, and benefits of AI innovation in government services: A review," *Discover Artificial Intelligence*, vol. 4, no. 1, p. 18, 2024. Available: <https://doi.org/10.1007/s44163-024-00111-w>
29. Blair Attard-Frost, Ana Brandusescu, and Kelly Lyons, "The governance of artificial intelligence in Canada: Findings and opportunities from a review of 84 AI governance initiatives," *Government Information Quarterly*, vol. 41, no. 2, p. 101929, 2024. Available: <https://doi.org/10.1016/j.giq.2024.101929>
30. A. Bhatnagar and V. R. K. Bonagiri, "Operational Excellence through Cloud-Native Distributed Systems," *International Journal of Computational and Experimental Science and Engineering*, vol. 12, no. 2, 2026. [Online]. Available: <https://doi.org/10.22399/ijcesen.5207>
31. H. Gaggar, "PromptDLP: Deterministic Pre-Inference Egress Control for Preventing Secret and PII Leakage in LLM Agents," *Journal of Information Systems Engineering and Management*, vol. 9, no. 1, 2024. [Online]. Available: https://www.jisem-journal.com/download/59_IJSEM_9_1.pdf
32. C. T. Bobba, "Advancements in Cosmetic Dentistry," *Innovation*, vol. 1, p. 2.
33. C. T. Bobba, R. Choksi, V. S. Grewal, H. K. Basra, and S. Momin, "Evaluation of Effects of Deproteinization and Intermediate Bonding on the Retention of Pit and Fissure Sealants-An In-Vivo Study," *SVOA Dentistry*, vol. 2, no. 4, pp. 139-142, 2021.
34. H. Gaggar, "Two Planes of Truth: Execution Fidelity and Behavioral Correctness in Tool-Using LLM Agents," *International Journal of Computational and Experimental Science and Engineering*, vol. 10, no. 3, 2024. [Online]. Available: <https://doi.org/10.22399/ijcesen.5498>