Volume 11, Issue 4, April – 2026 International Journal of Innovative Science and Research Technology
ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

Governing the Invisible Layer: AI Accountability Gaps Inside Enterprise Systems

Aditya Kashyap[1]

[1]Researcher in Enterprise Software Implementation,
AI Governance and Digital Transformation,
Bangalore, India

Publication Date: 2026/04/15

Abstract: AI governance frameworks such as the EU AI Act, the NIST AI Risk Management Framework, and ISO 42001 were built for a world where AI systems were visible, discrete, and clearly owned. That is not the world enterprises operate in today.AI now exists inside enterprise software. It influences decisions across ERP, CRM, contract management, and supply chain platforms. Organizations are not deploying AI systems. They are inheriting AI capabilities. This shift creates a structural accountability problem. Enterprises do not train these models and often cannot inspect them. Yet they remain fully accountable for the outcomes those systems produce. Vendors control the technology, but do not share operational responsibility. The result is a governance gap that current frameworks such as the EU AI Act, NIST AI RMF, and ISO 42001 are not designed to address. This paper explores how that gap appears in real enterprise environments and proposes the Enterprise AI Integration Governance model, a practical four-layer approach to reestablishing accountability in embedded AI ecosystems. With regulatory enforcement approaching, this is no longer a future concern.

Keywords: *AI Governance, Enterprise AI, EU AI Act, NIST AI RMF, ISO 42001, ERP, Embedded AI, Agentic AI, Accountability, Digital Transformation, AI Risk Management, CLM, SAP, Oracle.*

How to Cite: Aditya Kashyap (2026) Governing the Invisible Layer: AI Accountability Gaps Inside Enterprise Systems. *International Journal of Innovative Science and Research Technology*, 11(4), 676-683. https://doi.org/10.38124/ijisrt/26apr503

I. INTRODUCTION

Enterprise AI governance is often framed around the oversight of discrete, identifiable AI systems. In response, organizations are establishing ethics boards, developing responsible AI policies, and building compliance registers. These efforts are typically oriented toward AI systems that the enterprise can identify, scope, and govern directly.

This framing does not fully reflect how AI is introduced and operates within most enterprise environments.

For many organizations, AI capabilities do not arrive as standalone systems initiated through internal development or formal deployment processes. Instead, they are embedded within enterprise software platforms that already support core business operations. Capabilities such as demand forecasting in SAP S/4HANA, anomaly detection in Oracle Fusion Cloud, supplier risk insights in SAP Ariba, or autonomous workflow execution in Salesforce Agentforce, are increasingly delivered as part of vendor-managed product evolution.

In these contexts, AI functionality is introduced through software updates, configuration changes, or feature enablement rather than discrete project lifecycles. As a result, existing governance mechanisms such as change advisory boards, risk assessments, and model validation processes may not consistently recognize these activations as distinct AI deployment events requiring review.

This dynamic gives rise to a structural accountability challenge.

Enterprises remain responsible for business outcomes influenced by these systems, including operational decisions, financial reporting, and regulatory compliance. However, they typically do not control key elements of the AI lifecycle, including model development, training data selection, and ongoing model updates. While organizations may retain certain forms of operational control, such as configuration settings, input design, and output usage, they often lack transparency into model behavior and limited ability to independently validate system logic.

The result is a partial misalignment between accountability and control.

Current governance frameworks provide important foundations but are not fully optimized for this embedded AI context. The European Union Artificial Intelligence Act distinguishes between providers and deployers of AI systems, but does not explicitly address scenarios in which AI

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

capabilities are integrated into licensed enterprise software and activated without a clearly defined deployment decision. The NIST AI Risk Management Framework assumes that organizations have sufficient visibility and influence over the AI systems they govern, including the ability to map, measure, and manage risks across the lifecycle. ISO/IEC 42001:2023, as a management system standard, presumes that organizations can identify AI systems within scope and implement appropriate risk and impact assessment processes. These assumptions become more difficult to operationalize when AI capabilities are embedded, vendor-managed, and continuously evolving.

This gap is not theoretical. It reflects an increasingly common mode of AI adoption across enterprise software ecosystems and is likely to expand as vendors continue to integrate generative and agentic AI capabilities into core platforms.

This paper examines how existing governance frameworks operate under conditions of limited visibility and control, identifies specific failure modes in embedded AI environments, and proposes the Enterprise AI Integration Governance model as a practitioner-oriented approach for governing AI capabilities that are not directly developed or fully inspectable by the enterprise.

II. HOW AI IS EMBEDDED IN ENTERPRISE SYSTEMS TODAY

➢ *Three Deployment Classes:*

Embedded enterprise AI is not a single pattern. It operates across three distinct classes. Each class changes where decisions are made, how control is exercised, and what governance is required.

- **Copilot class AI** generates outputs for human use. These systems respond to prompts, summarize information, and draft content that supports decision making. SAP Joule, available across the SAP Business Technology Platform and expanding across S/4HANA, SuccessFactors, and Ariba, can respond to natural language queries, summarize HR records, draft procurement justifications, and generate financial narratives. Microsoft Copilot in Dynamics 365 performs similar functions across sales, service, and operations.

At a technical level, these systems sit as an interaction layer on top of enterprise data and application services. They combine large language models with retrieval mechanisms that pull structured and unstructured enterprise data into the response context.

The AI does not execute the decision. It shapes the information that precedes it. In practice, that distinction becomes less meaningful as usage scales. When outputs are generated continuously and review time is limited, the framing provided by the system strongly influences the outcome.

- **Decision assist class AI** produces scored, ranked, or flagged outputs that feed directly into business processes. Oracle Fusion Cloud uses machine learning models to evaluate journal entries during financial close and surface anomalies based on learned behavioral patterns. Contract analytics platforms extract obligations, identify non-standard clauses, and assign risk scores that influence whether agreements proceed, are escalated, or are renegotiated.

These systems are typically embedded within transaction processing layers. They operate on structured data, apply trained models, and write outputs back into workflow objects such as approvals, exceptions, or risk indicators.

Here, the AI output is not just advisory. It becomes the starting point of a process. While a human may remain in the loop, override conditions are often undefined, and the system output carries implicit authority.

- **Agentic class AI** executes sequences of actions across systems with conditional logic and limited human intervention. Salesforce Agentforce can manage service interactions by reading records, generating responses, updating case states, and triggering downstream workflows. Oracle has introduced AI agents in finance that can perform tasks such as journal matching, variance analysis, and elements of the period close process.

From an architectural perspective, these systems orchestrate multiple services. They combine reasoning capabilities with workflow engines, APIs, and event driven triggers to perform multi step operations.

In this class, the system is no longer shaping or initiating a decision. It is executing it. Actions occur at system speed and across system boundaries. Audit trails are generated as a byproduct of execution rather than as a designed control mechanism, which creates challenges for traceability and regulatory review.

➢ *The Scale of Exposure:*

This is not an edge case. It reflects the core of modern enterprise computing.

SAP software is used by the vast majority of the world's largest enterprises. Oracle Fusion Cloud ERP supports tens of thousands of organizations globally. Salesforce operates across a large and diverse enterprise and mid-market base. Each of these platforms is actively embedding AI capabilities across functional modules.

As a result, enterprises are not selectively adopting AI. They are inheriting it as part of the software they already run. These capabilities are introduced through updates, configurations, and feature enablement rather than discrete deployment decisions.

From an architectural standpoint, AI is becoming a horizontal layer across enterprise systems. It interacts with application logic, data models, and workflow engines without always being explicitly surfaced as a separate system.

The implication is straightforward. Every organization using these platforms is already operating AI within critical business processes. The difference is not adoption. It is awareness and governance maturity.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

III. STRUCTURAL LIMITATIONS OF EXISTING AI GOVERNANCE FRAMEWORKS IN EMBEDDED ENTERPRISE CONTEXTS

➢ *The EU AI Act: The Deployer Without Sufficient Visibility*

The EU AI Act, which came into force in August 2024, defines accountability through two primary roles. The provider builds and places an AI system on the market. The deployer uses that system in a professional context.

For most enterprises, this classification is straightforward. When an organization uses AI capabilities embedded in platforms such as SAP, Oracle, or Salesforce, it operates as the deployer. The obligations defined under Article 26 apply directly.

These obligations are clear. Enterprises must use systems in line with provider instructions, implement human oversight, monitor system behavior, and report serious incidents. In some cases, they must also perform data protection impact assessments under GDPR.

The challenge is not the existence of these obligations. It is the ability to fulfill them in practice.

Article 14 requires that human oversight be exercised by individuals who understand the system's capabilities, limitations, and outputs. This assumes that the deployer has meaningful visibility into how the system works.

In embedded AI environments, that assumption does not fully hold. From an architectural perspective, enterprise AI is not delivered as a standalone system with clear boundaries. It is integrated into application layers, workflow engines, and data services. The model operates behind APIs, orchestration layers, and vendor managed infrastructure. The enterprise interacts with outputs and configurations, not with the model itself.

This creates a visibility gap. The information required to support effective oversight, such as model behavior, training context, or known limitations, is only partially available through provider documentation. While vendors do supply instructions for use and certain transparency artifacts, these are not standardized across platforms and are often insufficient for deep interpretability at the level expected by Article 14.

As a result, the deployer is expected to exercise informed oversight over a system that is only partially observable. A similar tension appears in the classification of high risk systems. Annex III of the Act identifies use cases such as employment decisions and access to essential services as high risk. AI driven performance evaluation in systems like SAP SuccessFactors or financial decision support in Oracle Fusion may fall within this scope depending on how they are used.

In practice, the initial classification of whether an AI capability qualifies as high risk is typically driven by the provider's interpretation of the system's intended purpose. However, there is no consistent or standardized mechanism through which this classification is communicated to enterprise customers. The deployer must rely on limited disclosures while still carrying responsibility for compliant use.

This creates ambiguity at the point where classification drives obligation. The enforcement timeline further amplifies the issue. Prohibited practices became enforceable in early 2025. High risk obligations will apply from August 2026. This leaves a limited window for enterprises to establish governance mechanisms that align legal responsibility with operational reality.

The core issue is structural. The Act assumes a deployer that has sufficient visibility and control to govern an AI system. In embedded enterprise environments, the deployer has responsibility, partial control, and limited visibility. That gap is where governance begins to break down.

➢ *NIST AI RMF: Governing Without Sufficient Control*

The NIST AI Risk Management Framework organizes AI risk management into four functions. GOVERN defines accountability and policy structures. MAP establishes context and system understanding. MEASURE tracks risk through metrics and evaluation. MANAGE focuses on response and mitigation.

The framework is coherent. It assumes that the organization governing the system has meaningful visibility into how that system operates.

In embedded enterprise AI environments, that assumption becomes difficult to sustain.

GOVERN requires organizations to establish policies that apply across the AI lifecycle. In practice, embedded AI does not follow a lifecycle owned by the enterprise. It evolves through vendor-managed updates delivered via release cycles, configuration changes, and feature enablement.

From an architectural standpoint, the AI capability sits within vendor-controlled layers that combine models, data services, and application logic. The enterprise interacts with the system through interfaces and workflows, not through the underlying model components.

As a result, governance policies are often stable while the system they are meant to govern continues to change. The issue is not that policies become invalid, but that they are not tightly coupled to system evolution.

MAP requires organizations to understand how an AI system operates in order to identify risk. This includes documenting assumptions, limitations, and operational context.

For embedded AI, this understanding is indirect. The enterprise sees inputs and outputs. It may receive high-level documentation from the vendor. However, it does not have access to model architecture, training data, or detailed design decisions.

Architecturally, the model is abstracted behind service layers and APIs. This abstraction enables scalability and vendor control, but it limits the enterprise's ability to build a complete risk map at the level assumed by the framework.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

MEASURE requires that risks be tracked using defined metrics. This depends on observability. Organizations need to monitor outputs, detect performance changes, and identify emerging issues.

In embedded AI systems, observability is partial. Enterprises may have access to logs, usage patterns, and some output data. However, they typically do not have access to model-level telemetry such as distribution shifts, internal confidence signals, or degradation indicators.

This creates a measurement gap. Metrics can be defined, but they cannot always be supported by the data required to make them reliable.

GOVERN also recognizes the role of third-party AI. It calls for policies that address relationships with external providers.

This is directionally correct. However, the framework places the responsibility for establishing these policies on the enterprise, while the information required to operationalize them remains largely with the vendor.

The result is a structural imbalance. The enterprise is expected to govern risk across systems it does not design, does not fully observe, and does not directly control. The framework assumes a level of access and influence that does not consistently exist in embedded AI architectures.

That gap does not invalidate the framework. It defines the boundary of where it becomes difficult to apply without additional operating models.

ISO/IEC 42001: The Assessment That Becomes Difficult to Execute

ISO/IEC 42001:2023 is designed as a management system standard. Like ISO 27001, it focuses on establishing structure and discipline. Organizations are expected to define scope, understand their operating context, and implement processes to identify and manage AI risks.

At the center of this model is the assumption that AI systems can be identified, scoped, and assessed as they enter the organization.

In embedded AI environments, that assumption becomes difficult to sustain.

Clause 8.4 requires organizations to assess the potential impacts of AI systems as part of their lifecycle. This includes evaluating effects on individuals, groups, and broader outcomes, and defining appropriate controls.

For this to work, the organization needs a clear moment at which an AI system becomes part of its environment.

In practice, that moment is often unclear. AI capabilities in platforms such as SAP, Oracle, and Salesforce are introduced through product updates, configuration changes, and feature enablement. These changes are communicated through release notes or administrative settings, not through discrete deployment events that align with governance processes.

From an architectural perspective, the AI is not introduced as a separate system. It is integrated into application layers, data services, and workflow engines that already exist. The capability evolves within the system rather than entering it as something new.

This creates a gap between system change and governance recognition. The organization may have an impact assessment process, but it is not consistently triggered at the point where AI capabilities are activated. The assessment model assumes visibility into system boundaries. Embedded AI reduces those boundaries.

Annex A.10 addresses third-party relationships and requires organizations to define controls for vendor-provided AI. This includes setting expectations, monitoring compliance, and managing risk across external dependencies.

This is necessary, but not sufficient. Enterprises can define contractual requirements and governance expectations. However, the effectiveness of these controls depends on the information made available by the vendor. In many cases, disclosure is limited, non-standardized, or not aligned with the level of detail required for meaningful oversight.

Architecturally, this reflects a split in control. The vendor manages the model, infrastructure, and evolution of the system. The enterprise manages usage, context, and business outcomes.

The control model spans two parties, but the governance obligation sits primarily with one.

As a result, an organization can establish a well-structured AI management system, meet internal governance requirements, and still operate AI capabilities within core business systems that are only partially visible and not fully assessed.

The issue is not the absence of governance. It is the misalignment between where governance is applied and where AI actually operates.

➢ *Three Cases Where the Gap Produced Real Consequences*

- *Amazon Resume Screening, 2018.*

Amazon built an AI model to screen job applications, trained on a decade of historical hiring data. The model learned to systematically downgrade resumes from women, reflecting the gender imbalance in Amazon's own historical hiring outcomes. It operated inside an HR workflow — a decision-assist deployment — for approximately one year before the bias was identified through internal audit. Amazon discontinued the tool and did not disclose its use externally [13]. The governance failure was the absence of bias monitoring for an AI system embedded in a consequential HR process. Under EU AI Act Annex III point 4, this application would qualify as high-risk. Under NIST AI RMF MEASURE 2.5, bias metrics should have been defined before deployment. Neither mechanism existed.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

- *Zillow Offers, 2021.*

Zillow's automated home-buying business used an AI pricing model to generate purchase offers at scale. The model operated as an agentic system: it produced price recommendations that translated directly into binding purchase offers without individual human review at the transaction level. In Q3 2021, the model's predictions degraded significantly in a volatile housing market. Zillow recorded a $304 million inventory write-down in a single quarter and shut down the business unit, resulting in approximately 25% workforce reduction [14]. The model had been treated as reliable rather than probabilistic. No systematic mechanism existed to detect performance degradation and trigger human escalation before the damage accumulated. This is the absence of Layer 4 of the EAIG model — intervention architecture — applied at operational scale.

- *Moffatt v. Air Canada, British Columbia Civil Resolution Tribunal, 2024.*

A passenger sought a bereavement fare discount based on information provided by Air Canada's AI chatbot. The chatbot gave incorrect information about the bereavement policy. Air Canada's legal defense argued that the chatbot was a separate legal entity responsible for its own statements and that Air Canada could not be held liable for its outputs. The tribunal rejected this argument explicitly, ruling that Air Canada is responsible for all information on its website regardless of which system generates it, and awarded damages against the airline [15]. This ruling establishes a direct legal precedent: enterprises cannot disclaim liability for AI-generated outputs within their operational systems by attributing those outputs to the AI's autonomy. The accountability attaches to the enterprise, not the model.

IV. THE ENTERPRISE AI INTEGRATION GOVERNANCE (EAIG) MODEL

The EAIG model is not intended to replace existing frameworks such as the EU AI Act, the NIST AI Risk Management Framework, or ISO 42001.

Those frameworks define what responsible AI governance should achieve. They establish obligations, principles, and control expectations. The challenge lies in execution.

In embedded enterprise environments, organizations are expected to govern AI systems that they do not design, do not fully control, and cannot fully observe at the model level. The gap is not in guidance. It is in operationalization.

The EAIG model addresses this gap. It functions as an operating layer that translates regulatory and framework-level expectations into actions that enterprises can realistically execute within vendor-managed architectures. It aligns governance with the actual structure of modern enterprise systems, where AI capabilities are embedded within application layers, exposed through services, and evolved through vendor release cycles.

From an architectural perspective, EAIG sits between governance intent and system execution. It connects policy, risk, and compliance requirements to the points in the system where influence is still possible, such as configuration, workflow design, input control, and output validation.

The model is structured across four sequential layers. Each layer establishes the conditions required for the next. Without this progression, governance remains conceptual and cannot be applied consistently in environments where system boundaries are not clearly defined.

➢ *Layer 1: Classification*

Before an organization can define policy, assess risk, or assign compliance obligations, it must answer a more basic question. What AI is actually operating inside its environment.

In embedded enterprise systems, this is not a simple inventory exercise. AI capabilities are not introduced as standalone systems. They are activated through software updates, feature flags, and configuration settings within platforms that are already in use. Traditional tracking mechanisms that focus on approved AI projects do not capture this mode of deployment. As a result, a significant portion of enterprise AI can remain unaccounted for.

From an architectural perspective, these capabilities exist within application layers, workflow engines, and data services. They are integrated into existing processes rather than deployed as separate components. This makes them harder to isolate, identify, and classify using conventional system boundaries.

To address this, organizations need a structured AI census. This requires coordinated engagement with major software vendors, supported by vendor management and architecture functions. The objective is to establish a clear view of which AI capabilities are active within the licensed environment, which business domains they influence, how they function within workflows, and how they should be classified. This includes distinguishing between copilot, decision assist, and agentic patterns, as each carries different levels of risk and control.

It also requires understanding what documentation is available. This may include system descriptions, limitations, and any available governance artifacts provided by the vendor.

This is not a one-time activity. Because these capabilities evolve through vendor release cycles, the census must be maintained on a recurring basis aligned with those cycles. In practice, this often means reviewing changes quarterly or semi-annually, depending on the platform.

Regulatory classification adds a second layer of complexity. Under the EU AI Act, whether an AI system is considered high risk is determined based on its intended purpose and use case. Providers play a central role in this classification, particularly for systems placed on the market. However, enterprises as deployers must understand how these classifications apply within their own operational context.

For embedded AI features operating in critical domains such as finance, human resources, or customer decision making, organizations should explicitly seek classification clarity from vendors. Where this information is not clearly

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

provided, the enterprise operates with uncertainty around its regulatory obligations.

Classification, in this context, is not just labeling. It is the point at which visibility begins. Without it, governance cannot be reliably applied.

- *Layer 2: Accountability Assignment*

Most enterprise software contracts signed before 2023 did not anticipate embedded AI. They do not clearly define responsibility for AI generated outputs. They do not require vendors to notify customers before material changes to embedded AI capabilities. They do not allocate liability for decisions influenced by AI recommendations. In most cases, they do not reference emerging regulatory frameworks such as the EU AI Act.

This creates a contractual gap. AI capabilities are now embedded within the same systems that drive financial reporting, hiring decisions, procurement, and customer engagement. From an operational standpoint, many SaaS agreements now include AI driven functionality, whether explicitly recognized or not.

From an architectural perspective, the enterprise does not interact with the model directly. It interacts with application layers where AI outputs are already integrated into workflows and data objects. This means that accountability must be defined at the point where business decisions are made, not where the model is hosted.

To address this, enterprises need to update their contractual and vendor governance structures across four areas.

First, notification obligations. Vendors should provide advance notice for material changes to AI capabilities that affect consequential workflows. This includes model updates, changes in behavior, or expanded automation scope.

Second, liability clarity. Contracts should define how responsibility is allocated when AI outputs influence regulated decisions such as employment actions, financial reporting, procurement, or contract execution.

Third, regulatory alignment. Enterprises should require vendors to demonstrate compliance with applicable obligations, including transparency requirements under the EU AI Act, particularly where features may fall within high risk categories.

Fourth, documentation access. Enterprises should define expectations for access to relevant information such as system descriptions, known limitations, and available performance or bias assessments. The objective is not full model transparency, which may not be feasible, but sufficient information to support oversight and risk management.

The leverage for these changes exists. Regulatory obligations apply directly to providers under the EU AI Act. Enterprises operating in regulated environments can require evidence of compliance as part of contract renewal and vendor evaluation processes.

- *Layer 3: Observability*

Governance cannot function without visibility. For any AI influenced decision in a critical workflow, the organization should be able to reconstruct what happened. This includes the input provided to the system, the output generated, any associated scores or signals, the human action taken, and the resulting business outcome.

This expectation aligns with regulatory direction. The EU AI Act requires logging capabilities for high risk systems. The NIST AI RMF emphasizes the need for documentation that supports analysis and accountability.

From an architectural standpoint, this requires instrumentation at multiple layers. AI outputs must be captured within application data structures. Workflow systems must record how those outputs are used in decision processes. Audit and governance tools must integrate this information into a traceable record.

In many enterprise platforms, this capability does not exist by default in a complete form. It must be configured. This includes enabling audit logs, mapping AI generated fields within master and transactional data, and integrating these records into governance, risk, and compliance systems.

This work is not complex, but it is foundational. Without it, governance remains descriptive rather than operational.

- *Layer 4: Intervention Architecture*

Human oversight is a core requirement of modern AI governance. The challenge is not defining it, but implementing it meaningfully.

In many enterprise systems, oversight takes the form of a required approval step in front of an AI generated recommendation. At scale, this often becomes procedural rather than substantive.

From an architectural perspective, oversight must be designed into the workflow, not added as a checkpoint. Effective intervention depends on the class of AI being used.

For copilot class systems, the focus is on user capability. Individuals receiving AI generated outputs must have the domain expertise to evaluate them and the authority to reject or modify them.

For decision assist systems, the focus is on traceability. When AI generated scores or recommendations are not followed, the rationale should be recorded and periodically reviewed. This creates a feedback loop between human judgment and system behavior.

For agentic systems, the focus shifts to control mechanisms. Organizations need defined conditions under which automated actions are paused, escalated, or stopped. These circuit breakers must be tested in practice, not just defined in policy.

A useful indicator of oversight quality is how often human intervention occurs. Extremely low override rates may indicate that oversight has become procedural rather than

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

evaluative. This should be interpreted as a signal for review, not as evidence of system accuracy.

The objective is not to slow down systems. It is to ensure that control remains meaningful at the points where decisions carry real consequences.

V. OPERATIONAL IMPERATIVES FOR EMBEDDED ENTERPRISE AI GOVERNANCE

- *Run the Embedded AI Census First*

No other governance activity should begin before this. An organization may have well-documented AI policies, but if it has not identified which AI capabilities are active across its core platforms, it is governing an abstract model rather than its actual exposure.

From an architectural perspective, embedded AI sits inside application layers, workflows, and data services. It does not appear as a separate system. Without a structured census, these capabilities remain distributed and unaccounted for.

The census establishes the foundation. It identifies where AI is operating, how it is used, and how it should be classified. Every subsequent governance layer depends on this visibility.

- *Reclassify Vendor Contracts as AI Governance Instruments*

Many enterprise software agreements were not designed with embedded AI in mind. As a result, they do not clearly address responsibility for AI-driven outcomes, update notification expectations, or access to relevant system information.

In operational terms, these agreements now govern AI behavior within critical business processes.

Legal and procurement functions should review major vendor contracts against applicable regulatory obligations, including those under the EU AI Act. The objective is to incorporate provisions that address notification of material AI changes, allocation of responsibility, and access to sufficient documentation to support oversight.

From an architectural standpoint, control over AI is distributed between vendor-managed system layers and enterprise-controlled workflows. Contracts are the primary mechanism for defining how this shared control is governed.

- *Establish Observability Before Expanding Policy*

Governance requires evidence. Many organizations have developed AI policies, but the underlying systems do not consistently produce the records needed to enforce them. Without traceability, governance remains declarative.

For AI-influenced decisions, the organization should be able to reconstruct the sequence of events. This includes inputs, system outputs, any associated signals, human actions, and resulting outcomes.

Architecturally, this requires instrumentation across data, application, and workflow layers. Logging must capture AI-generated fields. Workflow systems must record how those outputs are used. Governance tools must integrate this information into auditable records.

This capability should be established before expanding policy frameworks. Without it, enforcement will remain limited.

- *Establish Governance Mechanisms Beyond IT Change Control*

Embedded AI does not follow traditional deployment models.

It is introduced through vendor-managed updates, configuration changes, and feature enablement. These changes often fall outside standard IT change control processes, which are designed around system modifications initiated within the enterprise.

As a result, AI capabilities can become operational without structured review. Organizations need a complementary governance mechanism that monitors vendor-driven changes. This includes reviewing release notes, identifying AI-related updates, routing material changes through targeted assessment, and documenting outcomes before those capabilities are widely used in critical workflows.

From an architectural perspective, this is a shift from project-based governance to continuous system monitoring.

VI. CONCLUSION

Enterprises are accountable for decisions influenced by AI, even when those systems are not internally developed or fully visible.

This is the practical condition created by embedded enterprise AI. Regulatory frameworks such as the EU AI Act define deployer obligations clearly. They do not fully address how those obligations are operationalized in environments where visibility and control are partial.

The EAIG model addresses this gap. Its four layers, classification, accountability assignment, observability, and intervention architecture, are not conceptual principles. They are operational conditions required to make governance frameworks implementable within embedded AI environments. The concept of governance debt is relevant here.

As with technical debt, it accumulates over time as systems evolve without corresponding governance controls. Each vendor release that introduces or modifies AI capabilities increases this exposure.

Unlike technical debt, the consequences are external. They surface through regulatory findings, legal challenges, and financial impact.

Organizations that demonstrate credible AI governance in the coming regulatory cycle will not be those with the most extensive policy documentation. They will be those that can answer four questions with evidence.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26apr503

- What AI is operating within the enterprise
- Who is accountable for its outputs
- Whether those outputs can be reconstructed
- Whether effective intervention is possible when outcomes are incorrect

Those answers define whether governance is functioning in practice.

REFERENCES

[1]. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonized rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L Series.
[2]. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce.
[3]. International Organization for Standardization. (2023). ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system. ISO.
[4]. SAP SE. (2023). SAP Joule: The generative AI copilot for business. SAP News Center.
[5]. Microsoft Corporation. (2024). Copilot in Dynamics 365 Sales overview. Microsoft Learn.
[6]. Oracle Corporation. (2024). Using AI and machine learning in Oracle Fusion Cloud Financials. Oracle Cloud Documentation.
[7]. Salesforce Inc. (2024). Introducing Agentforce: The agentic layer for the enterprise. Salesforce News. https://www.salesforce.com/news/press-releases/2024/09/12/agentforce-news/
[8]. Oracle Corporation. (2024). Oracle AI agents for finance Oracle Cloud World 2024. Oracle News.
[9]. SAP SE. (2024). SAP at a glance. SAP Investor Relations.
[10]. Oracle Corporation. (2024). Oracle Fusion Cloud ERP.
[11]. Salesforce Inc. (2024). Salesforce FY2024 Annual Report. Salesforce Investor Relations.
[12]. Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters.
[13]. Flint, J., & Mullin, B. (2021, November 2). Zillow quits home-flipping business, cites inability to forecast home prices. The Wall Street Journal.
[14]. British Columbia Civil Resolution Tribunal. (2024, February 14). Moffatt v. Air Canada, 2024 BCCRT 149.