Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Modernizing Procurement and Financial Controls Without Replacing Legacy ERP: A Governance-First AI Architecture

Vijaya Bhaskar Reddy Saadhu

Independent Researcher, USA

ARTICLE INFO

ABSTRACT

There is a fundamental architectural barrier to organizations in regulated industries using artificial intelligence for procurement and finance functions: regulation makes it extremely difficult, if not impossible, for organizations to change legacy enterprise resource planning systems. Governance frameworks in regulated industries demand accountability, traceability, auditability, and explainability of decision-making. Customary digital transformation approaches that assume the ability to adapt or replace the legacy systems are not fit for that purpose. We propose a governance-first reference architecture that applies non-intrusive external integration patterns, governed data planes, and AI service components to legacy ERP systems while treating auditability, traceability, explainability, and human oversight as cross-cutting architectural properties rather than as a compliance or ethical tier. The reference architecture does not consider governance only as a post-process step but rather as a design principle that governs all of the data flows, model lifecycles, and decision flows. Use cases in procurement and financial control, such as vendor risk scoring, invoice anomaly, and spend classification, illustrate the role of AI in risk detection, decision quality improvement, and process efficiency. They also show that AI does not replace decision-making concerning material matters: the architecture includes risk mitigation strategies against issues such as model drift, algorithmic bias, lack of model explainability, and data privacy. This phased approach to deployment is consistent with the enterprise's readiness for adoption and the concept of digital transformation as controlled augmentation. It provides regulated organizations with a realistic, compliance-friendly approach to developing and deploying the AI capabilities they need into an existing enterprise architecture without sacrificing control, regulatory defensibility, or institutional trust.

Keywords: Artificial Intelligence Governance, Legacy ERP Integration, Procurement Automation, Financial Controls, Explainable AI

1. INTRODUCTION

The adoption of AI across financial and procurement functions in enterprise is no longer an aspiration; it is becoming operationally imperative to tackle issues around fraud detection, spend visibility, vendor risk, and timely and accurate decision-making. For many regulated companies, these dreams run headlong into the fact that their legacy ERP systems are too deeply embedded within the business and validation landscape to be switched out or seriously customized [1].

While the literature on digital transformation holds that system modernization should happen before AI adoption, it can often be difficult to justify in heavily regulated environments, where the cost, risk, and disruption of replacing systems can be extreme. This ambition has often run into tension with the reality of existing systems, creating demand for non-intrusive approaches of extending AI functionality around existing systems.

In this paper, we introduce a governance-first reference architecture of a digital transformation of the procurement and financial controlling processes of legacy enterprise resource planning (ERP) systems, such that the governance (auditability, traceability, explainability, and human oversight) is a cross-cutting structural property rather than a

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

property featuring as a layer on top. Embedding governance thus provides the means by which AI can function in legacy-limited environments while still adhering to regulations and preserving control [3].

The article presents three contributions: (1) the design of a non-intrusive model for governing AI-driven transformation in constrained enterprise environments, (2) the derivation of a layered, governance-oriented reference architecture, and (3) an illustration of how AI can improve the quality of risk detection and decision-making without breaking the existing governance structures in procurement and financial control.

2. REGULATORY FRAMEWORK AND GOVERNANCE PRINCIPLES

Regulated industries rely on controls centered on financial reporting, data privacy, operational transparency, and decision accountability. The implementation of AI to augment procurement and financial controls introduces additional compliance risks that are not present in customary controls focused on human operators, including non-transparent algorithms, model drift, data leakage, and fairness violations [4].

In many jurisdictions, accounting standards require companies to have strong internal controls over financial reporting. This often means that information systems are responsible for producing records related to authorization, invoice approval, and fraud detection and documenting an unalterable audit trail for external auditors to review. [5] . The evidentiary burden is much greater if AI is used to make or influence material financial decisions.

Next, data privacy regulations may additionally require AI systems that access procurement data to comply with data minimization, purpose limitation, and data retention obligations if vendor financial data, contract terms, or contractor payments are considered personally identifiable information [6]. If procurement data is used to train the AI models, they should follow the principles of privacy by design, avoiding direct reconstruction of sensitive input from their outputs during inference.

Legal requirements are one aspect of industry-specific compliance. For example, financial services companies have to obey rules on operational resilience, and healthcare organizations have to comply with sector-specific privacy regulations in addition to the general data protection regulation [4].

Based on these imperatives, four principles can be derived. The first, accountability, requires ownership of decisions, documentation of the approval process, and a decision log. The second, Traceability, requires that every data element, every transformation, and every analytical step be logged and potentially be reconstructed. Auditing entails producing evidence that can be reviewed by independent auditors without the proprietary model parameters. Explainability considers that AI recommendations for high-stakes decisions need to be interpretable for decision-makers and auditable to regulators [3].

In the absence of such governance architecture, specific risks posed by AI systems (model drift in particular) can be exacerbated. Model drift occurs when training and test distributions are no longer the same, causing fraud detectors to overlook novel patterns. Algorithmic bias may systematically disadvantage certain vendor types or certain transaction types (i.e., compliance violations), creating discriminatory impact. For example, it is difficult to provide an explanation for black-box decision-making, creating regulatory liability [7]. The solution is governance that is built into the architecture rather than bolted on as an afterthought.

Regulatory Domain	Representative Framework	Key Governance Requirement	AI-Specific Risk Introduced
Financial Reporting & Internal Controls	SOX and equivalents	Immutable audit trails, segregation of duties	Unexplainable automated authorization decisions

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Data Protection & Privacy	GDPR, sector-specific rules	Data minimization, purpose limitation, retention compliance	Inference-based leakage of sensitive vendor data
Industry-Specific Mandates	Basel III, FERC, HIPAA	Operational resilience, sector control standards	Model drift undermining compliance monitoring
AI-Specific Regulation	EU AI Act	Transparency, human oversight, fairness, documentation	Algorithmic bias in vendor or transaction classification

Table 1: Regulatory Frameworks and Corresponding Governance Requirements for AI-Enabled Financial Controls [4][5][6]

3. LEGACY ERP BOUNDARIES AND NON-INTRUSIVE INTEGRATION ARCHITECTURE

The legacy characteristics of ERPs must be carefully analyzed when designing an integration strategy that accounts for legacy constraints while extending AI capabilities. Legacy characteristics of ERPs include their monolithic architecture, limited API surface area, batch processing nature, rigid data schema, and vendor lock-in [2]. These characteristics are not limitations; they are design drivers that should influence the way that every integration is designed.

The main limitation is that the ERP core is immutable and does not allow AI logic to be baked into ERP business processes or extend the ERP data model. All AI capabilities need to be implemented in external layers that interact with the ERP through approved contracts [1]. The integration must ensure that the ERP stays stable, available, and performant no matter what is going on outside of the ERP with its data.

Three patterns can address these constraints and vary by capability and cost. If an ERP vendor has published approved interfaces, then **API-based integration** is the best choice. These provide real-time extracts of ERP data. They are based on normalized adapter layers that convert non-homogeneous contracts between different API technologies into usable downstream data. For modern ERP software, **event-driven integration** via event stream-based APIs enables near real-time processing of business events, such as purchase order creation or invoice receipt. This makes it possible to trigger AI analyses when needed without having to rely on polling (see above) [8]. RPA may also be used as a last resort in situations where there are no publish/subscribe mechanisms available to extract data in bulk by simulating users interacting with the UI of ERPs with operational fragility and maintenance overhead.

All patterns use an integration design that fulfills three principles. The first principle is non-invasiveness: no changes to ERP code, database schema, or core business processes. Graceful degradation ensures that the malfunction of integrations does not affect the ERP itself. Audit readiness also requires that extracted data be logged with source, timestamp, volume, and status [3].

The data governance layer converts the raw data from the ERP system into a consumable format for the AI models. Data quality checks need to be performed continuously for completeness, accuracy, consistency, and timeliness in the ETL or ELT pipelines before moving the data to the data services for AI [9]. Tracking the lineage of data from the source through its transformations to the analytical result allows auditors to verify the sources of decisions. Implementing a single source of truth strategy avoids governance caveats where multiple systems would claim ownership over the same data entities, weakening clear accountability. The ERP is the authorizing source for

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

transactions, with the data layer extracting, validating, and staging the data, while other systems cannot override the ERP [9].

Integration Pattern	Availability Condition	Data Latency	Audit Capability	Operational Risk
API-Based Integration	Vendor-exposed approved interfaces	Real-time to near-real-time	High — discrete, loggable calls	Low — decoupled from ERP core
Event-Driven Integration	ERP event stream availability	Near-real-time	High — event logs by design	Low — non-invasive trigger model
Robotic Process Automation	No API or event stream available	Batch — scheduled extraction	Moderate — log-dependent on RPA design	Moderate — UI fragility on ERP changes
Batch File Export	Minimal integration capability	End-of-day or periodic	Low — limited granularity	Low operationally, high analytically

Table 2: Legacy ERP Integration Patterns and Governance Implications [2][8]

4. GOVERNANCE-ORIENTED REFERENCE ARCHITECTURE

The reference architecture recommends organizing the capabilities into six logical layers with governance dispersed throughout. It is compatible with the ERP legacy immutability and allows for transformation with AI for the procurement and financial abilities of the entire system.

Layer 1 (Legacy ERP Core) is unchanged. The ERP system remains the system of record, and the existing ERP controls (authorization workflows, SOD, and audit trails) are unchanged. **Layer 2 (Integration Layer)**: APIs & gateways, middleware, and adapters to extract data in a secure manner and transform it into the normalization of a data model for consumption. **Layer 3 (Data Layer):** Staged transformation, data quality, metadata & lineage management to make the governed data available for AI services. **Layer 4 (AI Services Layer)** includes machine learning models, inference engines, and analysis services, separated into distinct model-training environments and production inference environments. **Layer 5 (Control and Governance Layer)** implements cross-cutting components like policy engines, immutable audit logs, explainability features, and access control. **Layer 6 (User Interaction Layer)** includes dashboards, workflows, and alert interfaces where users can utilize AI recommendations while remaining in control of decision-making [3].

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

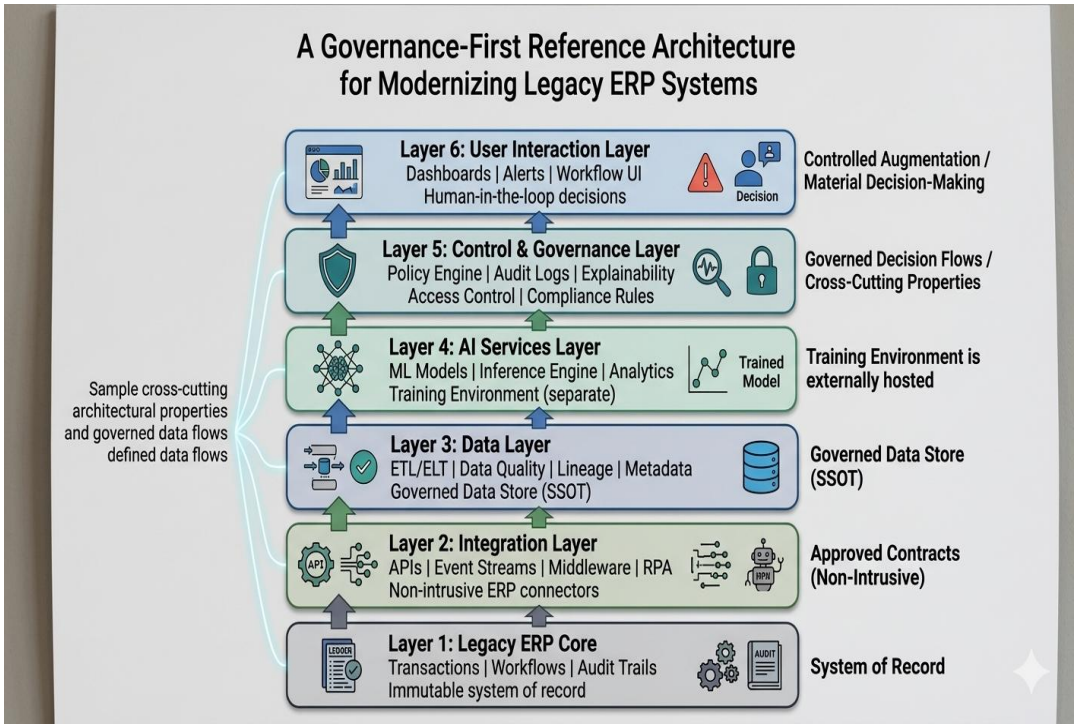

Fig. 1: Governance-Oriented Reference Architecture for AI-Enabled Transformation of Procurement and Financial Controls in Legacy ERP Systems

In this architecture, governance is a property of each layer rather than an independent Layer 5 function, and every extraction from the integration layer is traced and logged as such. The data layer documents every transformation and uses monitoring to identify poor data quality. The AI services layer uses model lifecycle governance. The user interaction layer logs every human decision, including rejection of recommendations made by AI [7].

The AI services layer is worth special mention since it includes the model lifecycle management capabilities. Models cannot be deployed to production without training and development, validation against held-out test datasets and historical backtesting, and the model has been through an AI governance board that reviews model documentation, its fairness and explainability, and its regulatory compliance. It is recommended to periodically monitor the model's performance and drift indicators after deployment and to revalidate the model every three or six months, depending on its criticality, to ensure continued performance under the current distribution of input data [10].

The policy engine encodes the rules of governance by associating checks with decision points. For example, vendor purchase orders from a high-risk vendor will be flagged for senior management approval, or invoices with a high anomaly score will be flagged for manual review prior to payment. These rules are configured rather than coded. Each time the policy engine receives an AI-generated recommendation, it checks it against the rules to ensure that the recommendation complies with the governance policies [3].

Human-in-the-loop decision-making is also a requirement for mission-critical decisions. Instead of making a decision to flag a fraud, block a vendor, or hold a payment, the AI systems would only provide a recommendation and a risk score. Humans review the output and explanation and have the option of approving, rejecting, or escalating. This creates a human-in-the-loop feedback process that holds the human responsible, catches false positives, allows feedback over time to improve the model, and ensures that organizations maintain control over critical decisions [11].

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

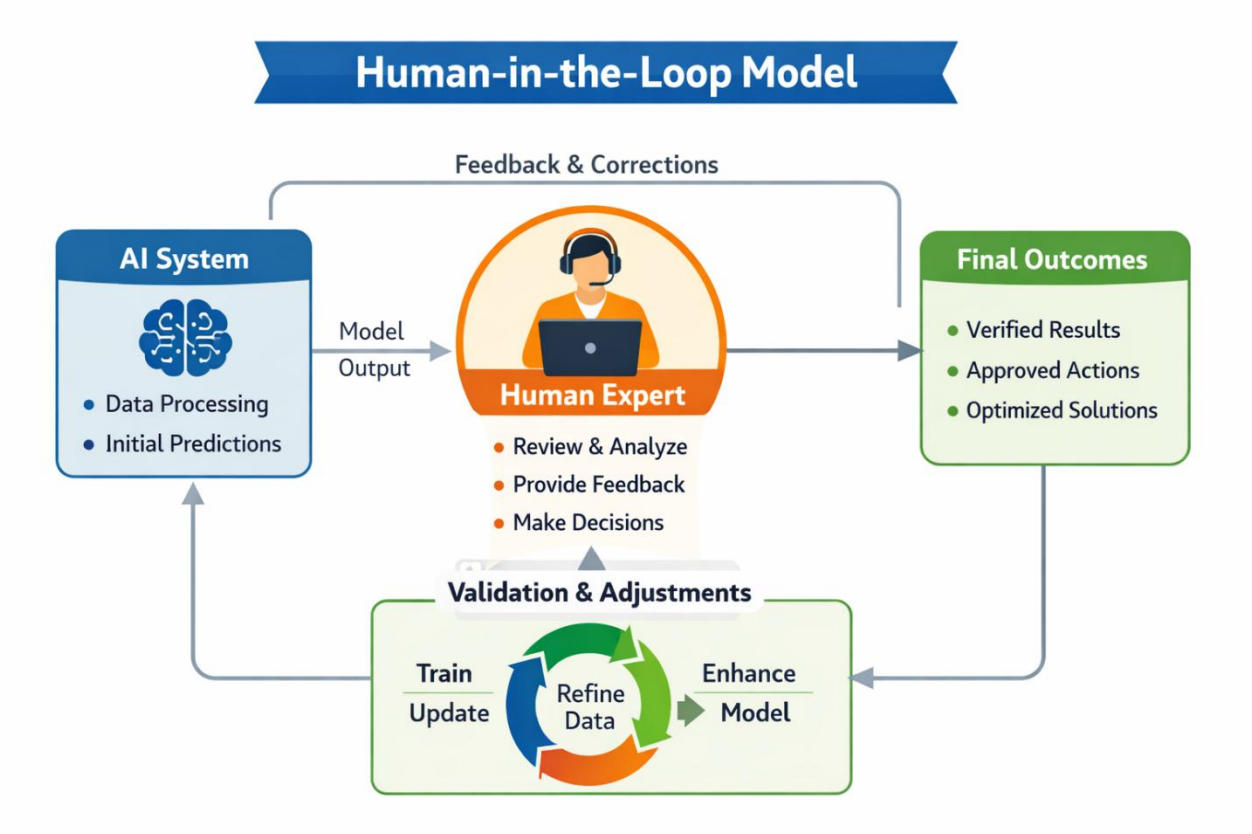

Fig. 2: Human-in-the-loop model

Architectural Layer	Primary Function	Embedded Governance Control	Audit Output Generated
Legacy ERP Core	Transactional system of record	Existing authorization workflows, segregation of duties	ERP-native transaction logs
Integration Layer	Data extraction and normalization	Extraction event logging, access control enforcement	Extraction of audit records with timestamps
Data Layer	Transformation, quality, and lineage	Data quality monitoring, lineage tracking, SSOT enforcement	Data quality metrics, lineage maps
AI Services Layer	Model training and inference	Model lifecycle governance, drift monitoring, versioning	Model approval records, performance dashboards
Control and Governance Layer	Policy enforcement and explainability	Policy engine rule execution, immutable decision logs	Decision audit trails, explainability reports
User Interaction Layer	Human review and decision authority	Human-in-the-loop mandate, override logging	User decision records with justifications

Table 3: Governance Architecture Layers and Embedded Controls [3][7]

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

5. PROCUREMENT AND FINANCIAL CONTROL USE CASES

Procurement Workflow Enhancement

One example is AI augmentation of existing processes. When a purchase requisition is created in the ERP, the integration layer pulls the requisition via API. The data layer will then improve the purchase requisition with historical spend against this vendor, the current vendor risk score, budget availability, and applicable contract terms . This also informs AI-driven vendor risk assessment and pricing anomaly detection [12].

The vendor risk scoring model calculates a risk score by standardizing various vendor factors into a payment history and regulatory compliance status, a financial stress measure, and a concentration risk measure. The spend analysis model compares a requisition cost of an item to historical unit costs from the same vendor. At runtime, the policy engine applies governance rules such as routing for senior approval if the vendor score is high risk, routing to a procurement manager if pricing is different from a threshold, or standard routing.

Human decision-makers receive the recommended, scored requisition and approve, deny, or defer it, including the timestamp, user, and justification. The approved requisition is sent back to the ERP as a PO through the integration layer with the added governance metadata. The entire process is tracked in the logs, which includes the requisition submission, AI analysis, policy routing, approval by the decision maker, and the creation of the PO [12].

Financial Control: Invoice Processing and Fraud Detection

AI can also help with high-volume invoice-based fraud. Multiple channels are used to submit invoices, and AI-enabled optical character recognition (OCR) can structure fields for downstream processes. Clean cases follow the standard three-way matching logic of invoice, purchase order, and goods receipt [13].

For those invoices, which could not be matched automatically, and those identified as anomalies by other means, fraud detection is conducted via a multidimensional approach. Duplicate detection is performed on invoice ID, vendor-amount-date combinations, and other fingerprints against the corpus of historical invoices. Amount anomaly detection checks invoice amounts against historical variations for each vendor. Vendor legitimacy checks validate vendor IDs in master files and check account age. Business logic violation checks are used to identify invoices with structural issues, such as post-dated transactions and mismatched line item totals.

The model's anomaly scores fuel the policy engine: high scores route invoices to fraud investigation teams, medium scores route them to accounts payable managers for review before payment, and low scores with clean three-way matches go straight to payment. Experts reviewing flagged invoices see the AI explanation alongside the invoice image and PO number ("flagged due to amount deviation, new vendor account age, invoice date inconsistency"), allowing informed decision-making rather than blind acceptance of AI.

Control objectives (e.g., segregation of duties, authorization, accuracy, timeliness, completeness, and accountability) are achieved by AI detection, policy-enforced routing, and human adjudication, with documentation for auditing purposes. This control model is best thought of as an augmentation to human control rather than a replacement: the AI detects anomalies at a scale which humans are unable to, and humans retain ultimate material decision-making authority [5].

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

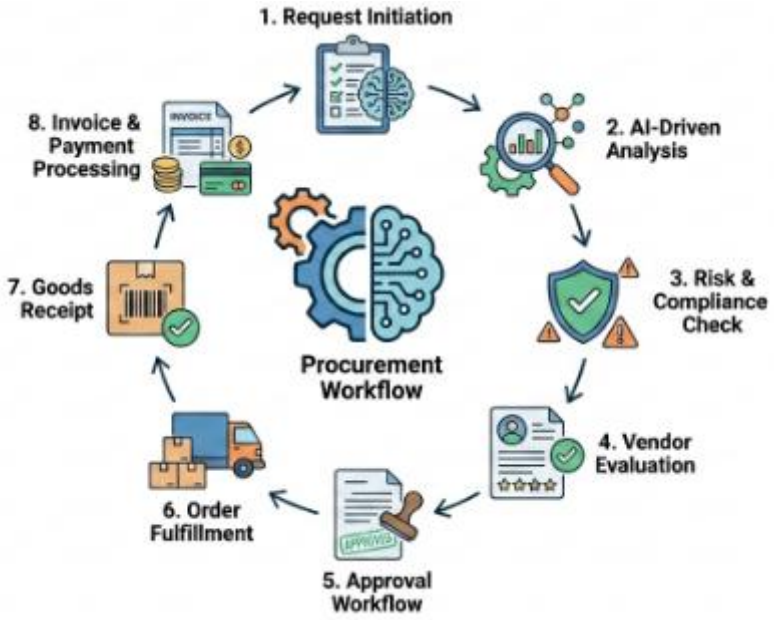

Fig. 3: Procurement workflow

Use Case	AI Capability Applied	Governance Control Applied	Human Decision Point	Control Objective Satisfied
Vendor Risk Scoring	Multi-factor risk model with explainability	Policy-driven approval routing by risk tier	Procurement manager approval for elevated risk	Authorization, accountability
Invoice Anomaly Detection	Anomaly scoring on invoice attributes	Escalation routing above anomaly score threshold	Finance specialist pre-payment review	Accuracy, fraud prevention
Pricing Anomaly Detection	Statistical deviation from historical unit costs	Procurement manager review trigger	Buyer review and justification	Accuracy, budget compliance
Spend Classification	AI-assisted GL and cost center coding	Data quality validation post-classification	Reconciliation review by finance	Completeness, spend visibility

Table 4: AI Use Cases in Procurement and Financial Controls with Risk and Governance Alignment [12][13]

6. RISK MANAGEMENT, IMPLEMENTATION PLAN, AND CONCLUSION

AI-Specific Risk Mitigation

It must also handle AI-specific risks, such as model drift (degradation of performance due to a shift in data distributions), by monitoring performance against baseline metrics, detecting statistical drift, automating retraining once thresholds are broken, and logging retraining events with triggers and resulting validation performance information [10]. Fraud detection models based on historic data must be regularly validated against the current fraud behavior to remain effective.

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Algorithmic bias creates compliance and ethical risk. Fairness metrics should be produced at both development time and production time to identify systematic differences in model predictions across vendor categories [7]. Further approaches include model explainability methods identifying the features used in the model's prediction and checking for inappropriate correlations with protected characteristics; for example, if the model uses vendor geography as an indicator of fraud risk, this feature will need to be removed and the model retrained. Other bias mitigation approaches include multiple approval processes for critical decisions.

To reduce black-box decision risk, the explainability-by-design model selection and mandatory explainability tooling are recommended. Shapley Additive Explanations (SHAP) estimate the contribution of each model input feature to each model output. They create evidence for decision-makers and auditors to understand and ensure the correctness of model decisions [14]. Local Interpretable Model-agnostic Explanations (LIME) are short, local approximations of bigger models. Regulatory documents describe how models work and are tested, and how explainability is ensured (kept to check) [14].

Data leakage and privacy are addressed through pre-deployment privacy impact assessments, feature selection, role-based access control of outputs, and compliance with data protection laws throughout the lifecycle of a dataset [6]. Overreliance on AI occurs when human reviewers no longer exercise independent judgment. This can be reduced through human-in-the-loop requirements, AI-override feedback, and periodic reviews of observed patterns of acceptance of decisions [11].

Phased Implementation

The first four phases of the deployment process, assessment and governance (months one through three), include defining the policies, the roles, and the decision-making processes. These must be established before the technical implementation begins. Phase two (months four to six) includes implementing the data governance layer, integrating the main sources, and enabling the governed staging environment, including quality monitoring and lineage tracking. Phase three (months seven to nine) should include implementing pilot AI use cases with a favorable profile for low operational risk and high learning potential, such as spend classification, invoice anomaly detection, and vendor risk scoring. These systems can be run in shadow mode (recommending actions without executing them) while awaiting human validation [3]. Finally, the fourth phase scales the pilots, instantiates vigilance dashboards, and defines thresholds for system escalation.

Risk Category	Definition	Architectural Mitigation	Monitoring Mechanism
Model Drift	Performance degradation from shifting data distributions	Automated retraining workflows triggered by drift thresholds	Continuous performance monitoring against baselines
Algorithmic Bias	Systematic disparate impact across vendor or transaction categories	Fairness metrics in development and production; feature auditing	Disparate impact ratio tracking in production
Black-Box Opacity	Inability to explain AI recommendations to decision-makers or auditors	SHAP and LIME explainability tooling is mandatory for material models	Explanation quality review during governance approval
Data Privacy Violation	Inference outputs are leaking sensitive vendor or contractor information	Privacy impact assessment, data minimization, output access control	Access log auditing, differential privacy, where applicable

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

AI Overreliance	Human reviewers deferring uncritically to AI recommendations	Human-in-the-loop mandate, override capture, periodic audit review	Decision acceptance rate monitoring with escalation thresholds

Table 5: AI Risk Categories, Mitigations, and Monitoring Mechanisms [10] [14] [15]

CONCLUSION

In existing regulated enterprise systems constrained by legacy, we have shown that governance-first augmentation (as opposed to replacement) is possible if governance is a structural architectural property rather than a compliance overlay. The governance-first reference architecture presented here shows that, with non-intrusive integration patterns, governed data pipelines, and externally hosted AI services, the legacy ERP core can be augmented with AI capabilities, while maintaining the integrity of existing transactional systems and established control models. For example, vendor risk scoring produces a multi-dimensional view of supplier risk that cannot be performed at scale, invoice anomaly detection detects fraud patterns that may not be detected through three-way match and consumer behavior and spend classification generates traceable visibility that can be used for budget compliance and expense governance. In all use cases human decision authority over material outcomes is retained, satisfying governance requirements in relation to both regulation and organizational accountability. The architecture supports risk mitigation against model drift, algorithmic bias, black-box model opacity, data leakage, and AI overreliance. By addressing these within the architecture, organizations can be defended over the long term as AI governance practices continue to evolve. The phased approach (from development of the governance framework to deploying data infrastructure and eventually implementing a pilot use case) allows for incremental adoption aligned with organizational readiness and risk profiles. As the global regulatory landscape for financial services and procurement technology matures, governance by design architectures represent the only scalable approach to institutionalizing AI in constrained enterprise environments and enable organizations to modernize mission-critical processes while maintaining the stability, compliance integrity, and institutional trust required for regulated operations.

REFERENCES

[1] Thomas H. Davenport and Rajeev Ronanki, "Artificial Intelligence for the Real World," Harvard Business Review, 2018. [Online]. Available: https://hbr.org/2018/01/artificial-intelligence-for-the-real-world

[2] Erik Brynjolfsson, Andrew McAfee, “The Second Machine Age." [Online]. Available: https://wwnorton.com/books/the-second-machine-age/

[3] Jeremy Pitt et al., "Algorithmic Reflexive Governance for Socio-Techno-Ecological Systems," IEEE Technology and Society Magazine, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9123456

[4] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 of the European Parliament and of the Council," Official Journal of the European Union, 2024. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

[5] Committee of Sponsoring Organizations of the Treadway Commission (COSO), “Internal Control — Integrated Framework," 2024. [Online]. Available: https://www.coso.org/guidance-on-ic

[6] European Parliament and Council of the European Union, “REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL," Official Journal of the European Union, 2016. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679

[7] Solon Barocas et al., “Fairness and Machine Learning: Limitations and Opportunities," 2023. [Online]. Available: https://fairmlbook.org/

[8] Martin Kleppmann, “Designing Data-Intensive Applications," Systems. O'Reilly, 2017. [Online]. Available: https://www.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/

[9] W. H. Inmon, Building the Data Warehouse, John Wiley & Sons, Inc., 2005. [Online]. Available: http://www.r-5.org/files/books/computers/databases/warehouses/W_H_Inmon-Building_the_Data_Warehouse-EN.pdf

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[10] João Gama et al., "A survey on concept drift adaptation," ACM Computing Surveys (CSUR), 2014. [Online]. Available: https://dl.acm.org/doi/10.1145/2523813
[11] Saleema Amershi et al., “Guidelines for Human-AI Interaction,” CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019. [Online]. Available: https://dl.acm.org/doi/10.1145/3290605.3300233
[12] Peter Kraljic, “Purchasing Must Become Supply Management,” HBR 1983. [Online]. Available: https://hbr.org/1983/09/purchasing-must-become-supply-management
[13] Richard J. Bolton and David J. Hand, "Statistical Fraud Detection: A Review," Project Euclid, 2002. [Online]. Available: https://projecteuclid.org/journals/statistical-science/volume-17/issue-3/Statistical-Fraud-Detection-A-Review/10.1214/ss/1042727940.full
[14] Scott Lundberg and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," arxiv logo > cs > arXiv:1705.07874, 2017. [Online]. Available: https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html
[15] Ninareh Mehrabi et al., "A Survey on Bias and Fairness in Machine Learning," ACM Computing Surveys (CSUR), 2021. [Online]. Available: https://dl.acm.org/doi/10.1145/3457607

Copyright © 2026 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.