iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VIII **Month of publication:** August 2026

DOI: https://doi.org/10.22214/ijraset.2026.84573

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VIII Aug 2026- Available at www.ijraset.com

Autonomous Decision Assurance Layer for AI-Driven Enterprise Analytics: Bridging Governance, Multi-Agent AI, and Executive Decision Intelligence in Saudi Vision 2030 Organizations

Choudhry Bilal Mazhar[1]
[#]*Digital Analytics & AI Lead, Arabian Agricultural Services Company (ARASCO), Riyadh, Saudi Arabia*

Abstract: ***Enterprise analytics is shifting from descriptive dashboards toward intelligent systems that recommend, prioritize, and increasingly trigger business actions. However, most organizations lack a structured assurance layer that validates whether AI-generated recommendations are explainable, risk-rated, auditable, policy-aligned, and suitable for execution. This paper proposes an Autonomous Decision Assurance Layer (ADAL) for AI-driven enterprise analytics environments. The proposed framework bridges data governance, multi-agent AI, human-in-the-loop oversight, responsible AI controls, and executive decision intelligence. ADAL introduces seven integrated components: data quality validation, AI recommendation generation, explainability mapping, risk scoring, human approval control, audit logging, and execution monitoring. The model is positioned for Saudi Vision 2030 organizations where digital transformation, governance maturity, cybersecurity readiness, and real-time decision intelligence are national and enterprise priorities. The paper contributes a practical governance-to-execution architecture that can be applied across IT service management, cybersecurity operations, workforce analytics, procurement anomaly detection, HSSE risk intelligence, and corporate performance management.***

Keywords: AI governance, decision assurance, multi-agent AI, enterprise analytics, human-in-the-loop, responsible AI, executive dashboards, Saudi Vision 2030

I. INTRODUCTION

Enterprise analytics has evolved from static reporting into intelligent decision-support environments that combine real-time data platforms, artificial intelligence, machine learning, workflow automation, and executive dashboards. While this evolution improves visibility and speed, it also introduces a new governance challenge: organizations must ensure that AI-generated recommendations are explainable, risk-aware, policy-compliant, auditable, and suitable for business execution before they influence operational or strategic decisions.

II. RESEARCH BACKGROUND AND MOTIVATION

Modern organizations increasingly rely on data lakehouses, semantic models, AI agents, and analytical dashboards to accelerate decision-making. However, the ability to generate recommendations does not automatically create trusted decision execution. A recommendation may be technically correct but operationally risky, insufficiently explained, based on poor-quality data, misaligned with policy, or unsuitable for automated execution. This paper addresses this gap by proposing an assurance layer that validates AI-driven recommendations before they become enterprise actions.

A. Research Gap

Existing enterprise analytics models usually focus on visualization, prediction, governance reporting, or automated workflow execution as separate capabilities. Limited attention has been given to the control layer between AI-generated insight and business action. This missing layer is especially important in regulated, critical, or high-impact environments where decision traceability, human oversight, and audit readiness are required.

This study contributes to the literature by defining ADAL as a practical control layer between AI-generated recommendations and enterprise execution. It extends existing AI governance discussions by translating responsible AI principles into operational assurance checkpoints that can be embedded within analytics platforms, workflow systems, and executive decision environments.

III. PROPOSED AUTONOMOUS DECISION ASSURANCE LAYER

The Autonomous Decision Assurance Layer (ADAL) is designed as a governance-to-execution control layer positioned between AI-driven analytics outputs and enterprise action workflows. Its purpose is to ensure that AI recommendations are not executed solely because they are technically generated, but because they meet defined standards for data quality, explainability, risk acceptability, operational relevance, and approval accountability. Prior research on AI governance, explainability, data governance, and human oversight provides important foundations for trustworthy enterprise AI. However, these streams are often treated separately: governance frameworks define principles, explainability research addresses model transparency, and enterprise analytics literature focuses on decision support. ADAL positions these areas within a single assurance layer that connects AI recommendations to controlled business execution. The paper adopts a conceptual design-science orientation, using governance requirements, responsible AI principles, and enterprise analytics practices to derive the ADAL framework. The framework is evaluated conceptually through applicability mapping, component alignment, and decision-assurance metrics rather than through empirical deployment.

TABLE I
CORE COMPONENTS OF THE PROPOSED ADAL FRAMEWORK

Component	Purpose	Assurance Output
Data Quality Validation	Checks completeness, freshness, consistency, and lineage of input data.	Quality score and readiness status
AI Recommendation Engine	Generates analytical recommendations using predictive, prescriptive, or multi-agent models.	Recommended action and confidence level
Explainability Mapping	Clarifies why a recommendation was produced and which variables influenced it.	Decision rationale and evidence trail
Risk Scoring	Assesses business, compliance, operational, cybersecurity, and reputational risk.	Low, medium, or high-risk category
Human Approval Control	Routes high-impact recommendations to authorized decision owners.	Approval, rejection, or escalation record
Audit Logging	Records decision inputs, model outputs, user actions, timestamps, and exceptions.	Audit-ready decision history
Execution Monitoring	Tracks whether approved actions produced the expected operational result.	Outcome monitoring and feedback loop

IV. REFERENCE ARCHITECTURE

The proposed architecture begins with enterprise data sources such as ERP systems, IT service management platforms, cybersecurity logs, HR systems, procurement systems, IoT sensors, and operational databases. These sources feed a governed data platform where quality, lineage, and access controls are applied. Analytical models and AI agents then generate recommendations, which are passed to the ADAL layer for validation before being displayed in executive dashboards or sent to workflow automation tools. The overall reference architecture is illustrated in Fig. 1.

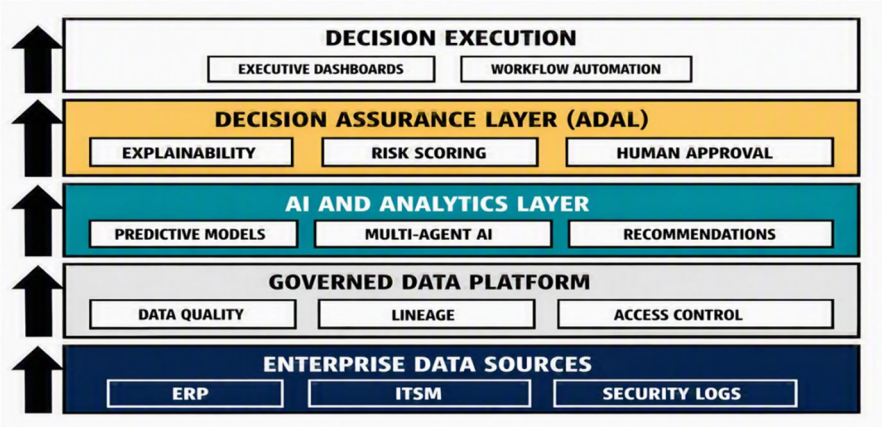

Fig. 1. Proposed ADAL reference architecture for AI-driven enterprise analytics.

As shown in Fig. 1, ADAL functions as the control point between intelligent analytics outputs and enterprise execution channels. The architecture supports progressive automation: low-risk decisions may be approved automatically when data quality, explainability, and policy thresholds are satisfied; medium-risk decisions may require operational owner review; and high-impact decisions require executive or governance committee approval. This tiered structure enables organizations to gain AI-enabled speed without compromising accountability.

V. DECISION ASSURANCE WORKFLOW

The ADAL workflow follows a structured sequence: data is validated, a recommendation is generated, explainability evidence is attached, risk is calculated, approval rules are applied, the decision is logged, and execution outcomes are monitored. Fig. 2 presents this workflow as a governed decision-assurance sequence that can be repeated across business domains while allowing each organization to configure thresholds based on its governance maturity and risk appetite.

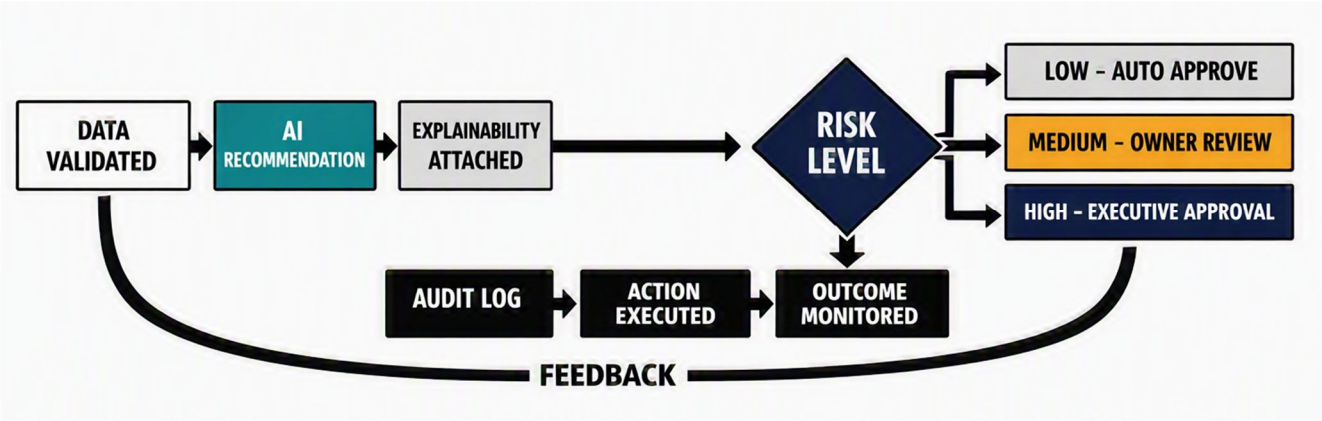

Fig. 2. Decision assurance workflow for governed AI recommendations.

Fig. 2 emphasizes that assurance is not a single validation event but a continuous governance cycle. Each recommendation is checked for input readiness, rationale transparency, risk exposure, approval requirements, and execution outcomes before it can influence enterprise action. This sequencing supports traceability from data input to final business result.

VI. USE CASE APPLICABILITY

The ADAL framework can be applied to multiple enterprise domains. Its value is highest where decisions are frequent, data-driven, operationally important, and subject to governance or compliance expectations. The following table summarizes representative use cases and the role of ADAL in each scenario.

TABLE II
ENTERPRISE USE CASE APPLICABILITY OF ADAL

Use Case	AI Recommendation	ADAL Assurance Role
IT Service Management	Prioritize incidents and assign support groups.	Validate severity, business impact, SLA risk, and approval path.
Cybersecurity Operations	Escalate suspicious events and trigger response actions.	Check alert confidence, asset criticality, and containment risk.
Workforce Analytics	Identify attrition, workload, or capability gaps.	Ensure ethical review, explainability, and privacy-aware decision controls.
Procurement Analytics	Detect spending anomalies or supplier risk.	Assess financial exposure, approval thresholds, and audit trail.
HSSE Risk Intelligence	Flag emerging operational safety risks.	Route high-risk cases for immediate human review and corrective action.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VIII Aug 2026- Available at www.ijraset.com

VII. EVALUATION CRITERIA

The effectiveness of ADAL can be evaluated using governance, operational, and technical metrics. These measures help organizations assess whether AI-enabled recommendations are reliable, explainable, and aligned with enterprise decision controls.

TABLE III
EVALUATION METRICS FOR DECISION ASSURANCE

Metric	Description	Expected Benefit
Data Readiness Score	Measures quality, completeness, and freshness of decision input data.	Improves trust in analytical recommendations.
Explainability Coverage	Measures the percentage of recommendations with clear rationale.	Improves transparency and user acceptance.
Risk Classification Accuracy	Assesses whether recommendations are assigned appropriate risk categories.	Reduces inappropriate automation.
Approval Cycle Time	Tracks time required to approve, reject, or escalate decisions.	Improves operational responsiveness.
Audit Completeness	Measures availability of logs, rationale, ownership, and outcomes.	Supports compliance and accountability.

VIII. DISCUSSION

The proposed ADAL framework supports responsible AI adoption by converting governance principles into operational control points. It is not limited to a specific vendor platform or industry domain; instead, it can be implemented through combinations of data platforms, dashboarding tools, AI services, workflow engines, and audit repositories. Its primary value is that it makes AI-powered decision-making measurable, explainable, and governable, reducing the risk that organizations move from insight generation to automated action without sufficient assurance.

For Saudi Vision 2030 organizations, ADAL offers a practical mechanism to strengthen digital transformation maturity. It supports faster executive decisions while maintaining human accountability, policy alignment, and audit readiness. This is particularly relevant for organizations operating in critical infrastructure, public services, smart cities, industrial operations, logistics, food security, and large enterprise environments.

IX. LIMITATIONS AND FUTURE WORK

This study proposes a conceptual and architectural framework rather than an empirically validated implementation. As a result, its effectiveness has not yet been tested through live enterprise deployment, controlled experiments, or longitudinal performance measurement. Future research should validate ADAL through case studies, prototype implementations, and quantitative evaluation across different organizational contexts. Further work may also examine sector-specific decision-risk models, automated policy mapping, integration with enterprise AI agents, and the role of synthetic data in safe decision simulation.

X. CONCLUSIONS

This paper introduced the Autonomous Decision Assurance Layer as a structured framework for bridging AI-driven enterprise analytics and governed business execution. The model addresses the gap between generating analytical recommendations and safely converting those recommendations into organizational actions. By combining data quality validation, explainability, risk scoring, human approval, audit logging, and execution monitoring, ADAL enables organizations to scale AI-assisted decision-making while maintaining trust, compliance, and operational accountability. The framework is especially relevant for Saudi Vision 2030 organizations seeking to advance analytics maturity, digital governance, autonomous operations, and executive decision intelligence in a responsible and auditable manner.

XI. ACKNOWLEDGMENT

The author acknowledges the broader digital transformation, analytics, responsible AI, and data governance communities whose work continues to shape enterprise adoption of trustworthy decision intelligence systems.

REFERENCES

[1] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg, MD, USA: NIST, 2023.
[2] ISO/IEC, Information Technology — Artificial Intelligence — Management System, ISO/IEC Standard 42001, 2023.
[3] European Parliament and Council, Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence. Brussels, Belgium: European Union, 2024.
[4] Organisation for Economic Co-operation and Development, OECD Principles on Artificial Intelligence. Paris, France: OECD, 2019.
[5] Saudi Data and Artificial Intelligence Authority, National Strategy for Data and Artificial Intelligence. Riyadh, Saudi Arabia: SDAIA, 2020.
[6] National Cybersecurity Authority, Essential Cybersecurity Controls (ECC-1:2018). Riyadh, Saudi Arabia: NCA, 2018.
[7] A. Batool, D. Zowghi, and M. Bano, "AI governance: a systematic literature review," AI and Ethics, vol. 5, pp. 3265–3279, 2025, doi: 10.1007/s43681-024-00653-w.
[8] A. Ghosh, A. Saini, and H. Barad, "Artificial intelligence in governance: recent trends, risks, challenges, innovative frameworks and future directions," AI & Society, vol. 40, no. 7, pp. 5685–5707, 2025, doi: 10.1007/s00146-025-02312-y.
[9] S. Hiremath, R. P., S. S. Konek, V. I. Bhavikatti, R. Vemula, and O. B., "Artificial intelligence (AI) governance in organizational decision-making: balancing autonomy, accountability and transparency," Journal of Entrepreneurship and Public Policy, 2025.
[10] "Explainable artificial intelligence in information systems: a review of the status quo and future research directions," Electronic Markets, vol. 33, 2023.
[11] M. Leon and H. DeSimone, "Advancements in explainable artificial intelligence for enhanced transparency and interpretability across business applications," Advances in Science, Technology and Engineering Systems Journal, vol. 9, no. 5, pp. 9–20, 2024.
[12] S. K. Maddala, "Understanding explainability in enterprise AI models," International Journal of Management Technology, vol. 12, no. 2, pp. 58–68, 2025.
[13] M. Balamurugan, K. Shanmugasamy, and S. Balaguru, "Humans in the loop, lives on the line: AI in high-risk decision making," European Journal of Computer Science and Information Technology, vol. 13, no. 51, pp. 27–31, 2025.
[14] D. Acev, S. Biyani, F. Rieder, T. T. Aldenhoff, M. Blazevic, D. M. Riehle, and M. A. Wimmer, "Systematic analysis of data governance frameworks and their relevance to data trusts," Management Review Quarterly, 2025, doi: 10.1007/s11301-025-00545-1.
[15] M. Maiti, P. Kayal, and A. Vujko, "A study on ethical implications of artificial intelligence adoption in business: challenges and best practices," Future Business Journal, vol. 11, no. 1, article 34, pp. 1–15, 2025, doi: 10.1186/s43093-025-00462-5.

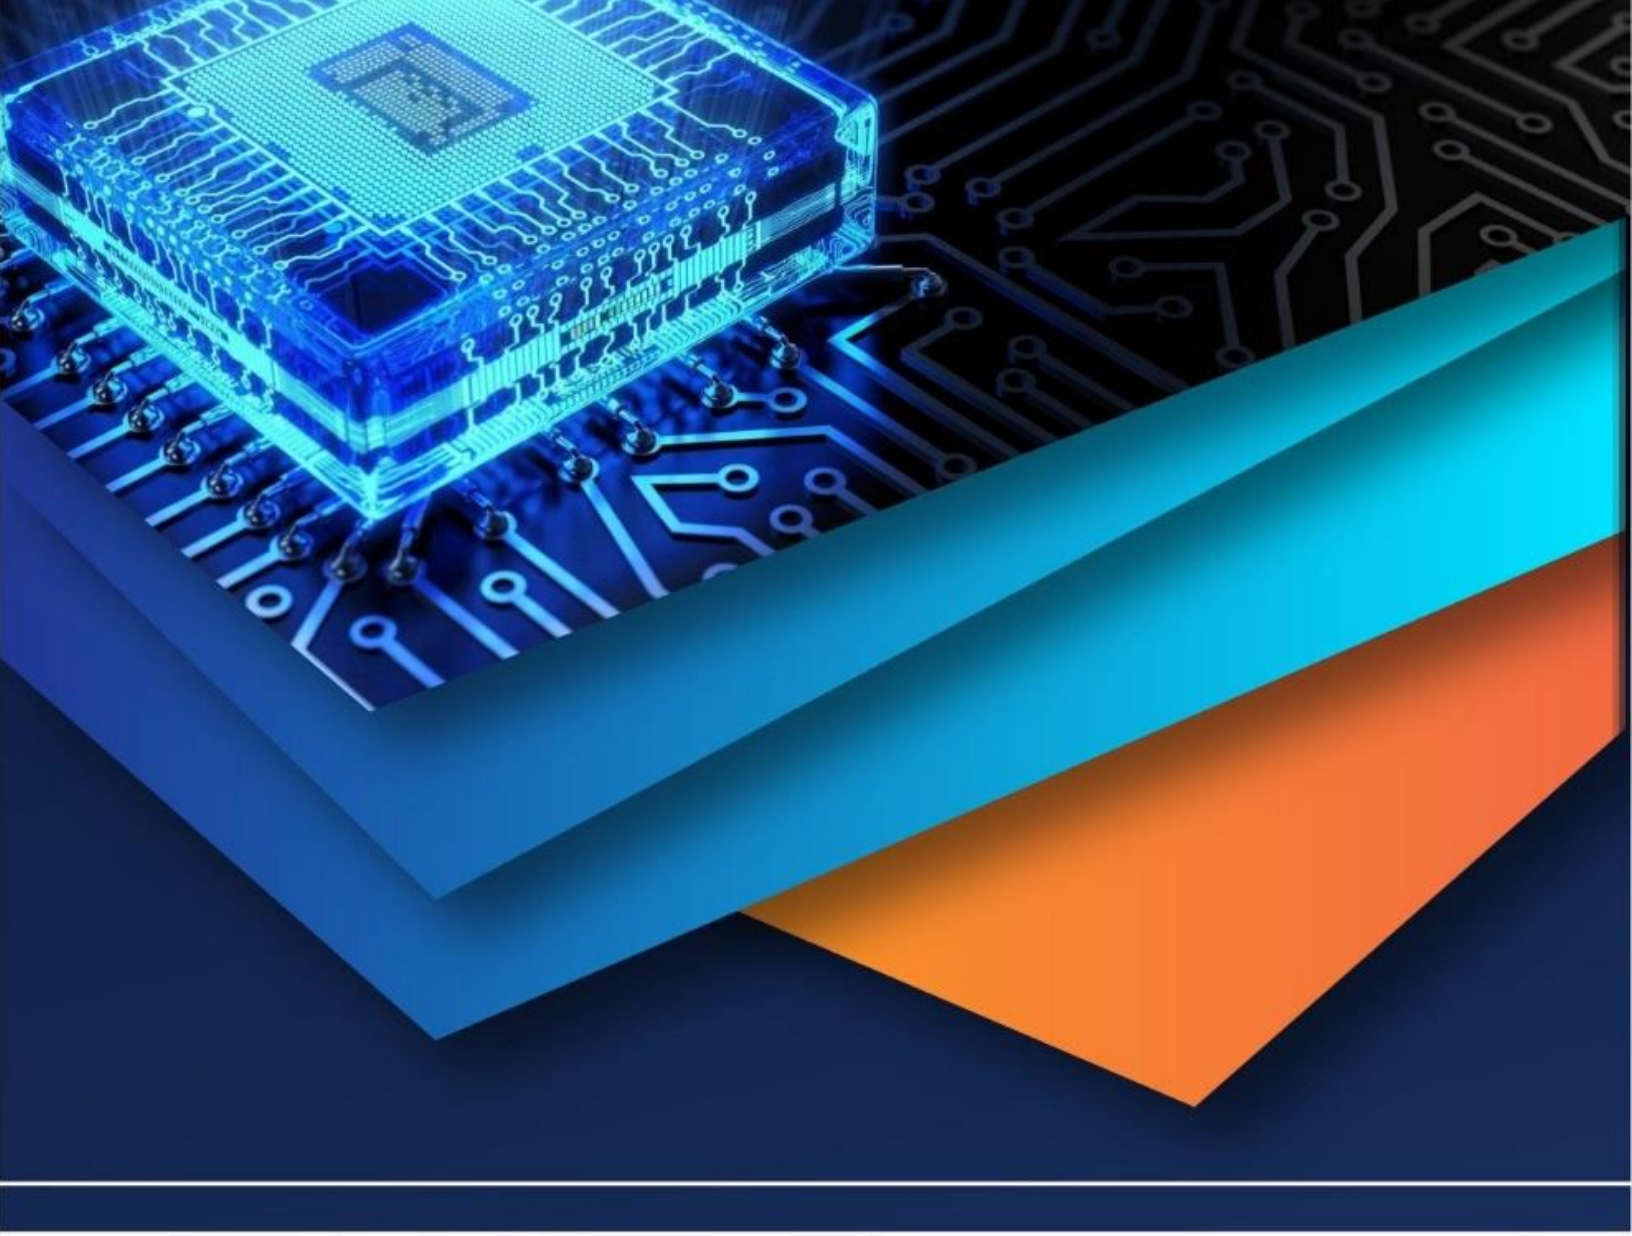

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY