© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722461

AI-Enabled Financial Risk Analytics: Reshaping Internal Audit and Enterprise Governance in Saudi Arabia

RIZWAN
CIA, CISA, AAIA, CFRA, CMA
Senior Group Internal Audit Manager, Riyadh, Kingdom of Saudi Arabia

Abstract- ***Artificial intelligence is moving from a peripheral analytics feature to the operating core of financial risk management in Saudi enterprises, embedded directly within SAP S/4HANA, Microsoft Copilot, Google's Gemini Enterprise, and purpose-built risk-scoring engines — and, increasingly, within the internal audit function's own tooling such as TeamMate+ and TeamMate Analytics. This adoption is not incidental: it follows directly from Saudi Arabia's Vision 2030 agenda, including the National Strategy for Data and AI, the Financial Sector Development Program, and the Human Capability Development Program, the last of which places a direct obligation on enterprises to build Saudi national capability alongside the technology itself. This creates a governance paradox for internal audit functions: the same AI capability that promises continuous, predictive assurance over credit risk, fraud, and control failure also introduces a new class of risk — model risk, algorithmic bias, and explainability gaps — that internal audit is not yet structurally equipped to test, including within the platforms auditors themselves rely on to produce evidence. Drawing on internal audit practice across diversified Saudi conglomerates spanning FMCG, QSR franchising, industrial distribution, energy, and real estate, this article examines how AI-enabled analytics are being applied to financial risk domains including expected credit loss estimation, fraud and anomaly detection, working capital forecasting, and continuous control monitoring. It maps these applications against the Kingdom's emerging AI governance architecture — SDAIA's Personal Data Protection Law (PDPL) and generative AI guidelines, the National Cybersecurity Authority's Essential Cybersecurity Controls, and international reference points including ISACA's AI governance guidance, COSO ERM (2017), and the IIA's emerging guidance on AI in internal audit. The article proposes a practical AI assurance model for internal audit functions, a structured approach to coaching Saudi talent into AI-enabled audit roles, and concludes with recommendations for enterprises seeking to adopt AI-enabled risk analytics without eroding the independence, objectivity, and evidentiary rigor that internal audit exists to provide.***

Keywords: Artificial intelligence, financial risk analytics, internal audit, enterprise governance, Saudi Arabia, Vision 2030, Human Capability Development Program, Saudi talent development, AI governance, algorithmic risk, model risk, SDAIA, PDPL, TeamMate+, Gemini Enterprise, continuous auditing, COSO ERM, machine learning, generative AI

I. INTRODUCTION

For most of its history, internal audit's relationship with technology has been reactive: systems change, and audit adapts its testing approach to keep pace. Artificial intelligence is breaking that pattern. Rather than sitting downstream of the financial process as another system to be tested, AI is increasingly embedded inside the process itself — scoring a customer's credit risk at the point of order entry, flagging a payment as anomalous before it clears, or drafting a variance commentary for management review. In Saudi enterprises undergoing simultaneous ERP modernization and Vision 2030-driven digital transformation, this embedding is happening faster than governance structures are adapting to it.

This creates a genuine paradox for the internal audit and risk function. AI-enabled analytics offer the most credible path yet toward continuous, full-population assurance over financial risk — something sample-based auditing has never been able to deliver. At the same time, AI models introduce a category of risk that traditional control testing was not designed to catch: a model can be operating exactly as configured and still produce a biased, poorly explainable, or silently degrading output. An audit function that simply extends its existing testing methodology to 'the AI system' without adapting its approach will miss the risks that matter most — and the same governance question now extends to the audit

function's own platforms, such as TeamMate+, as AI features are layered into risk assessment and workpaper preparation.

This article examines how AI-enabled financial risk analytics are actually being used inside Saudi enterprises today, the governance gaps this creates, and a practical model for how internal audit and enterprise governance structures can extend appropriate oversight without either rejecting the technology's genuine value or rubber-stamping its outputs.

II. VISION 2030 AS THE STRATEGIC CONTEXT

The pace at which AI is entering Saudi financial risk management is not incidental — it follows directly from priorities set under Vision 2030. Three elements of the national agenda bear most directly on how internal audit and enterprise governance should treat AI-enabled analytics.

2.1 The National Strategy for Data and AI (NSDAI)

The Kingdom's National Strategy for Data and AI sets an explicit ambition for Saudi Arabia to become a global leader in data and AI, with SDAIA established as the central regulatory and enabling body. For enterprises, this translates into both encouragement to adopt AI at scale and a rising expectation — reinforced through PDPL and SDAIA's generative AI guidance — that adoption be governed responsibly rather than left to organic, ungoverned uptake.

2.2 Financial Sector Development Program (FSDP)

The FSDP's push toward a more sophisticated, internationally benchmarked financial sector implicitly raises the bar for what audit committees and regulators consider adequate risk management technology. An enterprise still relying on static, backward-looking risk reporting increasingly reads as out of step with where the FSDP is steering the market, independent of whether AI adoption is formally mandated.

2.3 Human Capability Development Program (HCDP)

The HCDP is the Vision 2030 program most directly relevant to how enterprises should resource this transition. Its objective — building a Saudi workforce with the skills the future economy requires — applies with particular force to AI-enabled financial risk analytics, a discipline that barely existed as a defined skill set a few years ago and cannot simply be imported wholesale through expatriate hires without leaving a lasting capability gap in the national workforce. Section 8 addresses what this means in practice for how internal audit functions build capability, rather than treating it as a peripheral corporate social responsibility item.

III. WHERE AI IS ACTUALLY BEING USED IN FINANCIAL RISK TODAY

Across the enterprise portfolios reviewed in practice, AI-enabled analytics in financial risk management cluster around five use cases, each at a different level of maturity — four operating inside the business processes being audited, and a fifth inside the internal audit function's own tooling.

3.1 Predictive Credit and Receivables Risk Scoring

Machine-learning-based scoring models are increasingly used to supplement, and in some cases substitute for, traditional aging-bucket approaches to IFRS 9 Expected Credit Loss (ECL) estimation. These models incorporate payment history patterns, customer concentration, and macroeconomic indicators to produce a forward-looking probability of default, rather than the backward-looking aging bands that dominate manually maintained ECL spreadsheets. In practice, the accuracy gain is real, but so is the audit challenge: a model trained on historical payment behavior can embed exactly the biases and blind spots of the period it was trained on, and few finance teams can fully explain, without vendor support, why the model produced a specific customer's score.

3.2 Fraud and Anomaly Detection

Transaction-level anomaly detection — flagging unusual journal entries, duplicate vendor payments, or irregular expense claims — is the most mature AI use case observed, largely because it builds directly

on statistical outlier detection techniques that audit and forensic teams already understood before the 'AI' label was applied. The newer generation of tools adds pattern learning across large transaction populations, improving detection of subtler schemes such as gradual invoice splitting to stay under approval thresholds.

3.3 Working Capital and Liquidity Forecasting

Predictive forecasting of receivables collection, inventory turnover, and short-term liquidity position is being adopted particularly in sectors with volatile working capital cycles — QSR franchising and industrial distribution among them — where traditional static forecasting badly lags actual cash conversion behavior. These tools are generally viewed by finance teams as decision support rather than a control, which somewhat reduces the immediate governance urgency, but the same models are increasingly cited in board reporting as though they carried the same reliability as audited financial data.

3.4 Generative AI for Control Narratives, Risk Commentary, and Audit Documentation

Generative AI tools — including SAP Joule, Microsoft Copilot, Google's Gemini Enterprise, and general-purpose large language models — are being used to draft control descriptions, summarize risk register commentary, and in some cases produce first-draft audit finding narratives. Gemini Enterprise in particular is notable for the agent-building layer it adds on top of the underlying model: finance and risk teams can configure narrow-purpose agents against enterprise data sources rather than relying solely on a general chat interface, which raises the same governance stakes as SAP Joule and Copilot but with a wider surface area, since an enterprise may end up running several purpose-built agents rather than one sanctioned assistant. This use case carries the least direct financial risk exposure but the broadest governance blind spot, because generated text is fluent and plausible by design, which makes factual errors, fabricated figures, or subtly incorrect control descriptions harder to catch on review than they would be in a visibly rough manual draft.

3.5 AI Embedded in the Audit Function's Own Tooling

A distinct and often overlooked category is AI embedded not in the business processes being audited, but in the tools internal audit uses to perform the audit itself. TeamMate+, the audit management platform most widely used across the portfolio of Saudi enterprises reviewed in practice, has progressively layered AI-assisted capability onto its core risk assessment, planning, fieldwork, and issue-tracking workflow — alongside TeamMate Analytics, its standalone data analysis module used for full-population transaction testing. Where these features summarize workpaper content, suggest risk ratings, or draft finding language from fieldwork notes, they sit inside the audit evidence chain itself, which raises a governance question distinct from the ones in Sections 3.1 to 3.4: internal audit must be able to demonstrate that an AI-assisted risk rating or finding narrative reflects the auditor's professional judgment and the underlying evidence, not an unreviewed model output, if that workpaper is later relied upon by the audit committee, external auditors, or a regulator. This makes AI governance a quality-control matter for the audit function itself, not solely a subject the function tests in others.

IV. THE AI GOVERNANCE GAP IN INTERNAL AUDIT

Traditional internal audit methodology assumes that a control, once tested, behaves consistently until its configuration changes. AI models violate this assumption in three specific ways that most internal audit functions are not yet structured to address.

- Model drift: a model's accuracy can degrade silently over time as the underlying business or economic conditions diverge from its training data, with no configuration change to trigger a re-test.
- Explainability gaps: many AI models, particularly deep-learning-based ones, cannot produce a straightforward causal explanation for an individual output, which conflicts with internal audit's evidentiary standard of being able to trace a finding back to a specific, explainable cause.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722461

- Data lineage and bias: an AI model trained on historically biased data (for example, historical credit decisions that systematically disadvantaged certain customer segments) can reproduce and formalize that bias at scale, in a way that is far harder to detect than an individual biased human decision.

None of these gaps are unique to Saudi Arabia, but they intersect with a specific local regulatory layer — principally SDAIA's Personal Data Protection Law and its generative AI guidance — that gives them concrete compliance consequences rather than leaving them as a purely methodological concern.

V. THE SAUDI AI GOVERNANCE AND REGULATORY LAYER

Enterprises deploying AI-enabled financial risk analytics in the Kingdom operate against a regulatory backdrop that is still maturing but already carries real obligations.

Regulatory / Governance Reference	Scope	Practical Implication
SDAIA — Personal Data Protection Law (PDPL)	Governs collection, processing, and retention of personal data used to train or run risk-scoring models	Analytics teams must document lawful basis and data minimization wherever customer or employee personal data feeds an AI model
SDAIA — Generative AI Guidelines	Sets expectations for responsible and transparent use of generative AI tools	Generated content used in financial reporting or audit documentation should be clearly identifiable and subject to human review
National Cybersecurity Authority (NCA) — ECC	Establishes baseline cybersecurity control expectations, increasingly interpreted to cover AI system access and data handling	AI platforms and the data pipelines feeding them fall within the scope of existing cybersecurity control testing
CMA Corporate Governance Regulations	Requires boards of listed entities to demonstrate adequate oversight of material business risks	AI-enabled risk analytics used in board or audit committee reporting should be subject to explicit governance sign-off, not treated as a black-box input

Internationally, ISACA's AI governance guidance and the IIA's emerging position on AI in internal audit provide a complementary layer, framing AI oversight as an extension of existing IT governance and Three Lines Model principles rather than an entirely new discipline. The practical synthesis for a Saudi enterprise is a control taxonomy that treats AI governance as a cross-cutting layer over existing IT general controls, data governance, and enterprise risk management — not a standalone AI policy that sits apart from them.

VI. A PRACTICAL AI ASSURANCE MODEL FOR INTERNAL AUDIT

Based on this practice base, a four-layer assurance model is proposed for internal audit functions extending oversight into AI-enabled financial risk analytics.

Assurance Layer	Internal Audit Focus

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722461

Assurance Layer	Internal Audit Focus
Layer 1 — Data Governance	Validate the lawful basis, quality, and lineage of data feeding AI models, particularly where personal data is involved under PDPL.
Layer 2 — Model Governance	Confirm a named model owner, a defined validation and re-training cadence, and documented performance monitoring thresholds exist for each material model.
Layer 3 — Output Governance	Test whether AI-generated outputs used in financial reporting, board materials, or audit workpapers themselves (including TeamMate+ AI-assisted content) are reviewed, evidenced, and clearly distinguished from human-prepared content.
Layer 4 — Outcome Monitoring	Periodically re-test model outputs against actual outcomes (e.g., predicted versus actual default) to detect drift and bias rather than relying solely on point-in-time validation.

This model deliberately keeps internal audit in a testing and validation role rather than a model-building role. Where internal audit has been involved in configuring or tuning a risk model — a common temptation given the function's data fluency — independence must be preserved by separating that build activity from the team that subsequently provides assurance over the same model, consistent with IIA Standards on objectivity.

VII. ENTERPRISE GOVERNANCE IMPLICATIONS

Beyond internal audit's testing role, AI-enabled financial risk analytics raise questions that belong squarely with the board and audit committee.

- Board AI literacy: audit committees increasingly need enough technical fluency to ask whether a risk figure presented to them is model-generated, and if so, what validation stands behind it — a literacy gap that is itself an emerging governance risk.
- Accountability for model-driven decisions: where AI materially influences a credit, pricing, or provisioning decision, governance frameworks need to clearly assign accountability to a named business owner rather than allowing 'the model decided' to substitute for a decision-maker.
- Vendor and third-party AI risk: many AI capabilities arrive embedded in ERP or SaaS platforms rather than built in-house, meaning enterprise governance must extend to evaluating vendor model governance claims, not only internally built models.
- Escalation thresholds: enterprises benefit from defining in advance what magnitude of AI-flagged anomaly or risk score change triggers escalation to the risk committee, rather than leaving this to ad hoc judgment after the fact.

VIII. IMPLEMENTATION CHALLENGES

- Skills gap: internal audit teams generally have strong control-testing skills but limited data science fluency, making genuine model validation difficult without either upskilling or specialist support.
- Vendor opacity: enterprise software vendors do not always disclose model architecture or training data composition in sufficient detail to support independent validation, forcing audit functions to rely partly on vendor assurance reports.
- Documentation lag: AI model inventories, when they exist at all, are frequently incomplete or out of date relative to the models actually in production use, particularly where business units adopt AI features embedded in existing software without a formal onboarding process.
- Cultural resistance to challenge: because AI outputs carry an appearance of objectivity, business users and even auditors can be reluctant to challenge a model-generated figure the way

they would challenge a colleague's manual calculation — a bias internal audit training needs to explicitly address.

IX. COACHING AND DEVELOPING SAUDI TALENT FOR AI-ENABLED AUDIT

The Human Capability Development Program's ambition only becomes real inside individual audit functions if capability-building is treated as a structured, ongoing practice rather than a training budget line. In practice, four elements distinguish audit functions that are genuinely developing Saudi talent in this space from those that are merely staffing headcount targets.

9.1 Structured Rotational Exposure, Not Just Classroom Training

Data analytics and AI fluency are learned largely through supervised exposure to real engagements — reviewing an ECL model's outputs against actual receivables, validating a fraud-detection alert against underlying transactions, or shadowing the configuration of a GRC control monitor — rather than through standalone e-learning modules. Audit plans that deliberately rotate junior Saudi auditors through analytics-heavy engagements, paired with an experienced reviewer, build this fluency far faster than a certification course alone.

9.2 Certification Pathways Aligned to National Talent Goals

Structured progression through internationally recognized certifications — CIA, CISA, and increasingly AAIA for AI-specific audit competency — gives Saudi audit professionals a portable, benchmarked skill set that directly supports the HCDP's objective of a globally competitive national workforce. Audit functions benefit from formally sponsoring and tracking this progression as part of individual development plans rather than leaving certification pursuit to personal initiative alone.

9.3 Mentorship and Knowledge Transfer from Senior Practitioners

Where senior audit leadership includes practitioners with deep ERP, GRC, and AI governance experience, deliberate mentorship — co-reviewing workpapers, walking through the reasoning behind a risk rating, explaining why a model output was or was not accepted as sufficient audit evidence — transfers judgment in a way that documentation alone cannot. This is particularly important for AI-related audit judgment calls, which are new enough that formal methodology has not yet caught up with the nuance senior practitioners have already developed through hands-on exposure.

9.4 Building Internal Champions Rather Than Sole Dependency on Vendors

Vendor-led training on TeamMate+, SAP Analytics, or GRC modules builds initial familiarity, but sustainable capability requires developing one or two internal champions per platform — Saudi team members who progressively take ownership of configuration, troubleshooting, and onward training of colleagues — rather than the function remaining permanently dependent on vendor support desks or expatriate specialists for routine capability.

Taken together, these practices turn Vision 2030's Human Capability Development Program from a national-level ambition into a specific, measurable commitment inside the audit function's own people plan: every AI-enabled analytics capability the function adopts should have a corresponding development pathway for the Saudi talent who will own and advance it over time.

X. RECOMMENDATIONS

- Build and maintain a live AI model inventory covering every material AI application touching financial risk, with a named business owner for each — extending explicitly to AI features within the audit function's own platforms such as TeamMate+.
- Embed AI-specific test steps — data lineage, drift monitoring, explainability review — into the internal audit methodology rather than treating AI systems under generic IT general control testing.
- Require a PDPL data protection impact assessment before any AI model using personal data is deployed into financial risk processes, with internal audit periodically re-validating adherence.

- Invest in structured data analytics and AI literacy training for the internal audit function, building on existing certifications such as CISA and AAIA rather than treating AI competency as a separate specialism.
- Establish a standing AI governance item on the audit committee agenda, with clear escalation thresholds agreed in advance for AI-flagged financial risk anomalies.
- Formalize a Saudi talent development pathway alongside every new AI-enabled analytics capability the function adopts, consistent with the Human Capability Development Program, so platform ownership does not remain permanently dependent on expatriate specialists or vendor support.

XI. CONCLUSION

AI-enabled financial risk analytics represent a genuine step-change in what internal audit and enterprise governance can achieve in Saudi enterprises — full-population monitoring, forward-looking risk scoring, and faster detection of anomalies that manual sampling would likely miss entirely. But the same capability introduces a governance obligation that cannot be met by simply extending existing control-testing checklists to cover a new system category. Model risk, explainability, and algorithmic bias require a distinct assurance lens, grounded in data governance, model governance, and outcome monitoring, and reinforced by a Saudi regulatory architecture — PDPL, generative AI guidance, and cybersecurity controls — that is becoming more, not less, exacting. None of this is sustainable, however, if it remains dependent on a narrow pool of specialist expertise: Vision 2030's Human Capability Development Program makes clear that the same urgency applied to adopting AI technology must be applied to building the Saudi talent who will own, challenge, and advance it. Internal audit functions that build both the technical capability and the people pipeline now, rather than after a material AI-driven failure forces the issue, will be the ones enterprise boards can genuinely rely on as AI becomes the default rather than the exception in financial risk management.

REFERENCES AND SELECTED STANDARDS

[1] Kingdom of Saudi Arabia — Vision 2030, including the National Strategy for Data and AI (NSDAI), the Financial Sector Development Program (FSDP), and the Human Capability Development Program (HCDP).

[2] Institute of Internal Auditors (IIA) — International Professional Practices Framework (IPPF), the Three Lines Model, and emerging guidance on artificial intelligence in internal audit.

[3] ISACA — AI Governance, Risk and Auditing Tools guidance, and the AAIA (Advanced in AI Audit) certification body of knowledge.

[4] Committee of Sponsoring Organizations of the Treadway Commission (COSO) — Enterprise Risk Management: Integrating with Strategy and Performance (2017).

[5] International Organization for Standardization — ISO 31000:2018, Risk Management Guidelines; ISO/IEC 42001, AI Management Systems.

[6] International Financial Reporting Standards Foundation — IFRS 9 Financial Instruments (Expected Credit Loss requirements).

[7] Saudi Data and Artificial Intelligence Authority (SDAIA) — Personal Data Protection Law (PDPL), Implementing Regulations, and Generative AI Guidelines.

[8] National Cybersecurity Authority (NCA) — Essential Cybersecurity Controls (ECC).

[9] Capital Market Authority (CMA) — Corporate Governance Regulations, Kingdom of Saudi Arabia.

[10] Kingdom of Saudi Arabia — Vision 2030 and the National Strategy for Data and AI (NSDAI).

[11] SAP SE, Microsoft Corporation, Google Cloud — Product documentation for embedded AI features (SAP Joule, Microsoft Copilot, Gemini Enterprise) relevant to enterprise financial systems.

[12] Wolters Kluwer — TeamMate+ and TeamMate Analytics product documentation, audit management and data analytics tooling.