

Scaling AI readiness across fraud, identity, and money laundering mitigation

**Part 3: Operating-model considerations
for AI enablement**

This is a three-part series exploring the evolving financial crime landscape and the role of AI in helping organizations respond. In part 3, we focus on the operating model considerations needed to scale AI effectively, including integrating AI into investigations, governance, roles, training, change management, and human oversight.

Payments organizations have entered a period where fraud, identity, and money laundering risks are becoming increasingly interconnected. AI is accelerating the speed, scale, and sophistication of financial crime, enabling bad actors to industrialize tactics such as synthetic identities, deepfakes, phishing, and account takeover.

In response, institutions are moving beyond isolated automation efforts toward a broader redesign of data, models, investigations, governance, and operations. For payments leaders, this is not just a technology modernization agenda; it is a strategic imperative to scale AI readiness in ways that improve detection, strengthen decision-making, and support more effective, enterprise-wide mitigation of converging risks.



Operating-model considerations for AI enablement

As financial institutions expand the use of AI, the primary constraint is often not the technology itself but the operating model surrounding it. The institutions that progress furthest will not necessarily be those with the most sophisticated models, but those that align structure, incentives, governance, talent, and performance management around a more integrated view of financial crime risk.

Fraud, identity, and money laundering are adjacent, but not interchangeable

Fraud and money laundering increasingly intersect, but they remain distinct functions with different objectives, service levels, and accountability structures. Fraud is typically optimized for speed, intervention, customer protection, and loss mitigation, while money laundering is oriented toward investigative judgment, suspicious activity identification, documentation quality, and regulatory defensibility.

The practical design principle is to treat fraud and anti-money laundering (AML) as adjacent financial crime functions with shared infrastructure, while preserving separate scorecards, decision rights, and governance aligned to each function's distinct objectives. This allows AI to improve detection, triage, and investigative effectiveness without creating fragmented ownership, conflicting incentives, or faster decisions that do not lead to better outcomes.

There is also a practical issue around handoffs. As fraud, AML, and identity-related risks continue to converge, institutions need documented criteria for when a matter stays within one team and when it becomes a shared case. Without that clarity, organizations create inconsistent referrals, duplicated review, and fragmented evidence. Institutions therefore need defined escalation triggers, shared case-linking standards, and clear ownership across the life cycle of a financial crime event.



The challenge of quantifying AML outcomes from AI

A central organizational challenge is that AML success is inherently harder to quantify than fraud success. Fraud programs can often point to relatively direct metrics, such as losses avoided, false positive reduction, recovery rates, or investigation cycle time. AML programs, by contrast, are typically evaluated through proxy indicators such as alert quality, suspicious activity report (SAR) productivity, timeliness, model performance, and examination outcomes.

The relevant question is not whether AI increases activity, but whether it helps the institution make better decisions with greater consistency and appropriate human oversight. A more useful scorecard combines efficiency and effectiveness. Efficiency metrics may include cycle time, backlog stability, duplicate review reduction, handoff quality, and the percentage of low-value manual work removed from the process. Effectiveness metrics should focus on typology coverage, quality of escalation decisions, consistency of outcomes, cross-team referral value, case-linkage rates, and the stability of model performance over time.

Roles, training, and change-management considerations

In many institutions, AI education is focused first on technical teams, but the larger gap is often among operational users and control functions. Investigators need to understand what a model score means, what it does not mean, and when it should be challenged. Managers need to distinguish between prioritization support and automated decisioning. Oversight teams need enough familiarity with AI methods to review them credibly. And senior leaders need to recognize that governance cannot be added after deployment. Without that shared literacy, organizations tend to swing between two equally problematic outcomes: over trusting the technology or underusing it.

Change management is equally important. Resistance to AI in fraud and AML environments is rarely just about the tools themselves; it usually reflects concern about accountability, judgment, and workload. Investigators may worry that AI-generated summaries oversimplify complex cases. Compliance teams may worry that AI-supported recommendations weaken defensibility. Operations leaders may worry that added governance slows urgent action. Those concerns should be addressed directly through phased deployment, transparent operating rules, clear override standards, and closed-loop feedback between users, builders, and oversight teams. Adoption improves when institutions define exactly where AI adds value and where human judgment remains non-negotiable.



Integrating AI into investigations

Integrating AI into fraud and AML investigations is no longer just an automation opportunity. It is increasingly a broader redesign of investigative work. Europol's 2025 SOCTA report notes that nearly all serious and organized crime now has a digital footprint, meaning investigation teams are managing larger data volumes, more cross-channel evidence, and faster-moving activity than traditional manual case models were designed to handle.¹ AI can improve the consistency of case documentation, the completeness of evidence packaging, and the prioritization of investigator time toward higher-risk matters that require human judgment.

Europol's 2025 SOCTA report notes that nearly all serious and organized crime now has a digital footprint

Large language models (LLMs)

LLMs are most effective in investigations when deployed as workflow accelerators rather than stand-alone decision-makers. In investigations and operations, the highest-value use cases include summarizing alerts and case history, extracting facts from unstructured documents, drafting investigation narratives, identifying missing information, and recommending next actions for investigator review. These capabilities are especially relevant

to case triage, evidence synthesis, and SAR-supported activities, where teams must move efficiently through large case volumes while maintaining clear documentation and defensible reasoning. However, final escalation, disposition, and filing decisions should remain human-led.




Agentic AI

Agentic AI pushes the model from “assist” to “act.” Unlike a standard LLM assistant, agentic AI can perform multistep workflows with limited human prompting. For example, gathering data across systems, generating an investigation plan, pulling supporting records, and routing work to the appropriate queue. In investigations, this creates a meaningful step-change because it enables automation of end-to-end alert handling rather than simply accelerating isolated text-based tasks.

This is especially relevant in sanctions operations, where agentic AI can automate first-level adjudication of alerts by collecting and assessing match data, comparing it against customer and transaction context, and generating a preliminary disposition with supporting rationale. Some reports suggest that automating these steps can reduce SAR drafting time by roughly 70%.² At the same time, human-in-the-loop remains essential as AI-generated content needs to be validated for accuracy, and the accountability for final SAR decisions continues to rest with human investigators and compliance leaders.

Figure 1. Agentic AI in multistep investigations



Risk management and oversight

A practical model is therefore tiered automation. Low-risk, repetitive activities such as alert enrichment, documentation gathering, summarization, and first-pass recommendations are the strongest candidates for AI-led execution. Higher-risk judgments, particularly those involving escalation decisions, regulatory interpretation, customer impact, or formal filing accountability, should remain human-led. This operating model allows firms to scale investigative throughput without transferring regulatory responsibility away from accountable personnel.

At the same time, implementation needs to be paired with disciplined risk management. NIST’s Generative AI Profile, issued in July 2024, identifies the major risk categories that are directly relevant in investigative settings: confabulation or false outputs;

harmful or misleading content; privacy and data leakage risks; bias and information integrity issues; intellectual property concerns; misuse for cyber or disinformation purposes; excessive agency where systems take unauthorized action; and supply chain vulnerabilities involving models, data sets, or application programming interfaces (APIs).

For financial institutions, these risks are not just technology risks; they also create clear control requirements.³ Deploying LLMs and agentic AI in investigations therefore requires clear guardrails around human review, access controls, auditability, testing, monitoring, and role definition so that AI can accelerate investigative work without weakening governance, defensibility, or regulatory accountability.⁴

How Deloitte can support your organization

Scaling AI across fraud and AML requires deep experience in business strategy, data, technology, operating model, and governance so that AI can be deployed effectively, responsibly, and at enterprise scale. We can help your team:

- **Define the AI operating model across fraud, identity, and AML:** Establish how AI should be deployed across adjacent fraud, identity, and AML functions, including the business objectives, governance model, decision rights, and accountability structure needed to support sustainable adoption.
- **Strengthen identity, biometrics, and cyber defenses against AI-enabled fraud:** Drawing on Deloitte's Cyber Identity, Cyber AI, and Digital Identity+ capabilities, we can help you modernize onboarding, authentication, and ongoing customer monitoring—leveraging physical and behavioral biometrics, liveness and deepfake detection, and phishing-resistant authentication.
- **Secure frontier and agentic AI deployments end to end:** Apply Deloitte's Secure Software Development Lifecycle (SSDL), Managed Extended Detection and Response (MXDR), and

AI-aware threat-modeling capabilities to harden LLMs and autonomous agents, both internally and externally. We can help increase your defense against prompt injection, memory poisoning, tool misuse, model supply chain attacks, and excessive agency with controls extending from intent and guardrails through to runtime monitoring, red teaming, and post-deployment remediation.

- **Design a layered detection architecture:** Develop a detection stack that brings together rules, scorecards, machine learning, network analytics, and LLM-enabled workflow support. This also includes defining a platform strategy that balances enterprise consistency and vendor capabilities with institution-specific models, features, and decisioning logic.
- **Embed AI into investigative workflows:** Apply AI to triage, case summarization, evidence development, and investigator support activities to help reduce manual effort and improve consistency.

Contact us

Satish Lalchand

Principal
Deloitte Transactions and Business Analytics LLP
slalchand@deloitte.com

Anish Srivastava

Principal, Cyber Leader for Financial Services
Deloitte & Touche LLP
anissrivastava@deloitte.com

Nikhil Aggarwal

Managing Director
Deloitte Transactions and Business Analytics LLP
nikhaggarwal@deloitte.com

Brendan Maggiore

Senior Manager
Deloitte Transactions and Business Analytics LLP
bmaggiore@deloitte.com

Jonathan Thomas

Senior Manager
Deloitte Transactions and Business Analytics LLP
jthomas@deloitte.com

Endnotes

1. Europol, [*European Union Serious and Organised Crime Threat Assessment – The changing DNA of serious and organised crime*](#), Publications Office of the European Union, 2025.
2. FinCrime Intelligence, [“AI-driven SAR drafting in financial crime compliance: Real cases, real possibilities,”](#) accessed July 2026.
3. National Institute of Standards and Technology (NIST), [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\) \(NIST AI 100-1\)](#), January 2023.
4. NIST, [Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile \(NIST AI 600-1\)](#), July 2024.



As used in this publication, 'Deloitte' means Deloitte & Touche LLP, which provides audit, assurance, and risk and financial advisory services; and Deloitte Transactions and Business Analytics LLP, which provides eDiscovery, analytics, and a wide range of advisory and analytic services. These entities are separate subsidiaries of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.