

Deloitte.

Together makes progress

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

Part 2: Architecting AI for detection



This is a three-part series exploring the evolving financial crime landscape and the role of AI in helping organizations respond. In part 2, we look at how institutions can architect AI for better financial crime detection by combining stronger signals, more connected data, and more effective model design.

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.



Moving from a point solution to a layered capability

Better financial crime detection won't come from adding one more model to an already fragmented control stack. It requires a detection architecture that brings together identity, device, behavioral, transaction, and investigative signals into a shared analytical layer. For financial institutions, that means detection architecture should

be designed as a layered capability rather than a point solution. Institutions therefore need architectures that combine stronger identity proofing, liveness and anti-spoofing controls, cross-channel entity resolution, and AI models that can continuously learn from fraud, identity, and anti-money laundering (AML) outcomes together.


Common features and model characteristics

Models leveraged across fraud and financial crime include common features and model characteristics. To enable scalable fraud and AML detection, institutions need infrastructure that brings together the full range of signals used across both domains into a shared analytical layer. An effective pipeline should unify internal operational data, customer and account information, transaction histories, authentication and device signals, case management records, and external risk intelligence; fragmented data environments make it harder to detect cross-channel financial crime patterns and slower to develop reusable models. Strong programs combine structured inputs such as transactions, know your customer (KYC) data, beneficial ownership records, sanctions information, device and network data, and behavioral history with adverse media, case notes, call records, and other investigative documents.

A useful way to operationalize this defense is to anchor signals in a Root of Trust that is established through sequential stages of identity proofing, resolution, validation, and verification, that move from probabilistic to deterministic confidence, consistent with NIST SP 800-63 digital identity guidance.¹

- **Resolution** distinguishes a unique identity in the context of the system using probabilistic signals such as geolocation, network checks, behavioral and bot detection, session tracking, and device data that flags AI-driven scripted account creation and emulator-based attacks.
- **Validation** collects identity attributes including name, date of birth, address, SSN/EIN, phone, and email and confirms their accuracy and internal consistency against authoritative sources, surfacing the fabricated-data combinations that synthetic identities rely on.
- **Verification** then deterministically binds the validated identity to the live person through government-issued photo ID, face matching, facial liveness with passive and active anti-spoofing, age estimation, anti-fraud database checks, and synthetic identity checks.

Figure 1: Identity proofing stages for establishing Root of Trust
Resolution, validation, and verification, moving from probabilistic to deterministic confidence



Many leading institutions are pairing physical biometrics with liveness signals, such as behavioral biometrics, that continuously authenticate users through keystroke dynamics, mouse movement, device handling, and navigation patterns that are difficult for synthetic actors to replicate.

These signals strengthen onboarding, step-up authentication, and transaction monitoring—helping to link seemingly unrelated synthetic accounts that share devices, networks, or behavioral fingerprints, turning an onboarding control into a cross-domain financial crime capability.

Large language models (LLMs) for feature extraction

LLMs are increasingly valuable because they can convert unstructured content, including transaction narratives, adverse media, and investigator notes, into structured features such as risk topics, entity mentions, jurisdiction flags, behavioral indicators, and typology classifications. For example, adverse media can be converted into structured features including the frequency, recency, and severity of negative mentions; the types of alleged misconduct cited; and the strength of association between a customer and high-risk entities, industries, or geographies.

Third-party data sources

Third party data sources improve entity resolution, onboarding risk assessment, and transaction monitoring by helping institutions identify risk signals and relationship patterns that would not be visible from internal data alone. For example, sanctions, politically exposed persons, and watchlist data can be converted into features such as match confidence, proximity to listed parties, jurisdiction

risk, and the number of linked high-risk entities. Similarly, bureau and consortium fraud data can generate features such as prior identity misuse, chargeback or dispute history, account tenure across institutions, velocity of new-account openings, and reuse of devices, phone numbers, emails, or bank accounts across unrelated parties.

Applying network analysis for feature development

As criminal activity becomes more interconnected across individuals, entities, accounts, and devices, financial institutions need frameworks that can detect these relationships and the shared infrastructure supporting coordinated activity. By analyzing the relationships across customers, accounts, devices, counterparties, and transactions, institutions can generate features that capture patterns of connection, coordination, and behavioral similarity across the broader ecosystem. This helps identify suspicious activity such as mule networks, circular fund flows, smurfing behavior, shared identifiers, and other relational patterns that traditional linear monitoring often misses.

For example, network analysis can reveal that several seemingly unrelated accounts share devices, phone numbers, addresses, or counterparties and move funds in coordinated patterns. These connections can be translated into features such as shared-identifier counts, transaction cluster density, and proximity to previously suspicious entities to improve downstream fraud detection models.



Using entity resolution and segmentation to reduce false positives upstream

A scalable detection architecture improves alert quality upstream by strengthening the features and analytical context used before alerts are generated, rather than relying primarily on downstream alert suppression. Entity resolution supports this by improving the accuracy and completeness of the risk context at the point of detection.

By linking records associated with the same customer, business, account, device, or related party across fragmented systems, it can help institutions distinguish legitimate connected activity from unexplained or potentially suspicious relationships. For example, multiple accounts sending funds to the same beneficiary may initially appear indicative of funneling or structuring, but entity resolution

may determine that the accounts belong to related legal entities within the same corporate group.

Segmentation can also help reduce false positives by assessing activity against the appropriate behavioral baseline, recognizing that patterns that appear anomalous in one customer segment may be routine in another.

For example, transaction velocity, cash activity, or cross-border payments that may appear suspicious for a retail customer may be expected for a multinational commercial client. These capabilities can help distinguish genuinely suspicious

Building custom models for specific typologies

That same logic supports the case for custom models by typology rather than relying solely on one monolithic enterprise model. A synthetic identity model, mule account model, transaction laundering model, or crypto-related typology model may each require different labels, features, thresholds, and escalation paths because the underlying data, behaviors, and red flags are materially different. In many mature programs, those typology-specific models

sit beneath a broader ensemble layer that combines model scores, rule hits, network indicators, and investigator feedback into a higher-level prioritization framework. Scorecards are especially useful in this architecture because they offer a transparent and governable way to combine multiple risk indicators into a composite score, making them effective as baseline models, challenger models, onboarding tools, or prioritization layers within a broader AI ecosystem.

Figure 2: Illustrative scorecard approach

Risk Indicator	Observed	Points	Notes
Recent account opening	✓	15	New account opened within 30 days
High transaction velocity	✓	20	Multiple rapid transactions in short period
Shared device across accounts	✗	10	No shared device detected
Rapid cash-out behavior	✓	25	Funds withdrawn quickly after deposit
Prior alert history	✗	15	No prior investigator alerts
Crypto exposure	✓	10	Activity linked to crypto exchange
Composite Risk Score = 70		70	

0-29 Low 30-59 Medium 60+ High Priority Review

Platform strategy: Vendor solutions versus in-house development

For many institutions, the question is not a binary build-versus-buy choice, but what to buy as a core platform and where to build. Building internally offers control and potential differentiation, especially where the institution has proprietary data, distinctive risk patterns, and strong engineering and governance maturity. But it is also often costly, slower to deploy, and harder to sustain once validation, model monitoring, documentation, and regulatory change are factored into maintenance.

Vendor platforms can accelerate deployment, provide specialized data and AI capabilities, and reduce the operational burden of maintaining the full stack, which is why many institutions adopt a hybrid model: buy the core workflow and infrastructure, then layer in custom typology models, internal features, or proprietary decisioning where differentiation matters most.

Hosting strategy: Cloud, on-premise, and hybrid deployment


The same principle applies to hosting strategy. Model hosting should be treated not only as a technology choice, but also as an architectural and governance decision shaped by integration feasibility, inference speed, model oversight, and third-party risk. Many institutions are adopting hybrid approaches, using cloud environments for scalable compute and advanced AI capabilities while retaining certain data, systems, or models on-premise where legacy architecture, latency requirements, control considerations, or

data sensitivity make that more practical. This is especially relevant where core source systems remain on-premise but must connect to cloud-based AI models through application programming interfaces (APIs), middleware, or orchestration layers without disrupting transaction processing or case workflows.

An illustrative AI-enhanced architecture

AI-enhanced architecture for fraud, identity, and AML detection integrates source data, enriched and transformed data, monitored models, and feedback loops into a unified detection framework.

Figure 3: Illustrative AI-enhanced fraud, identity, and AML detection architecture



How Deloitte can support your organization

Scaling AI across fraud and AML requires deep experience across 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.

In part 3, we will focus on the operating model considerations needed to scale AI effectively, including integrating AI into investigations, governance, roles, training, change management, and human oversight.

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. National Institute of Standards and Technology (NIST), [Digital Identity Guidelines: Identity Proofing and Enrollment \(NIST SP 800-63A, 2nd public draft\)](#), 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.