Deloitte.

Together makes progress



Seeing the whole attack

Fraud intelligence for the AI era

The challenge

Fraud actors are escalating their operational speed beyond the capacity of existing preventative programs. The proliferation of large language models (LLMs) enables malicious actors to rapidly develop novel, high-impact fraudulent schemes. The emergence of "Fraud-as-a-Service" platforms monetizes successful attack methodologies, making these sophisticated methods available even to less-skilled fraud actors. Agentic AI further diminishes the cost associated with initiating a higher volume of attacks with increased frequency. Established fraud tactics are being repackaged and deployed, outpacing the adaptive capabilities of conventional rule-based detection systems.

Many fraud prevention programs are challenged by the lack of a standardized definition of fraud, making it difficult to establish a view of the threat across the organization. This is compounded by organizational silos that fraud actors recognize and exploit. While individual domains like cyber, transactions, operations, identity, or third-party activities may have robust detection capabilities, these efforts often remain focused on identifying fraud within each domain rather than understanding how fraud activity connects across them. As a result, organizations have limited insight into how a fraud actor's tools, tactics, and practices are evolving and how those changes manifest across the enterprise. Shifting away from a primarily detection-focused approach toward an intelligence-led mindset can help organizations connect signals, anticipate emerging threats, and identify opportunities to intervene earlier.

Our point of view

Deloitte's approach starts by providing counter-fraud teams with a unified, structural way to define the adversary. This Fraud Actor Framework delineates the perpetrator's identity, objectives, attack vectors, and operational methodologies, with the motivations (why) and timing (when).

From there, we focus on the fraud intelligence lifecycle: a continuous process of collecting, processing, analyzing, and connecting signals across domains; translating those signals into actionable intelligence; and using the resulting insights to inform and improve the next cycle. The lifecycle is informed by the fraud kill chain (the sequence of activities an adversary undertakes to execute a scheme, from target selection and reconnaissance, through attack execution, persistence, and disengagement or re-engagement). Intelligence generated at any point of the kill chain can interdict the attack, with earlier identification generally enabling a faster, more effective, and less costly response. By connecting signals across domains, organizations can move beyond siloed defenses and develop a broad view of the adversary and the threat. The intelligence lifecycle powers our Fraud Action Operating Model: the day-to-day structure that puts the intelligence into motion.



How it works: three phases

01

Understand & collect

Each engagement starts with a detailed mapping of fraud actor profiles against the organization's operational landscape. This involves identifying involved parties (both internal and external); defining their objectives (e.g., monetary assets, data, property, services); pinpointing the most vulnerable attack surfaces; and predicting how the attack will likely unfold across human, process, and technological vectors. Upon establishing this mapping, a governed ingestion layer is deployed to assimilate transactional, payment, logistics, cyber, open-source intelligence (OSINT), geospatial, and human-derived signals. This data is acquired via batch processing, application programming interface (API) integrations, and real-time streaming feeds, with schema validation and entity resolution applied at the point of entry to facilitate consistent entity linkage from inception.

02

Process & predict

Raw signals require standardization, connectivity, and broad modeling to yield actionable intelligence. These signals are consolidated into a unified data hub, with confirmed fraudulent events incorporated into a purpose-built fraud actor database. Subsequently, advanced analytics are applied, leveraging established classical models such as gradient boosting and random forests for real-time transaction scoring, recurrent and convolutional networks for identifying behavioral patterns, and clustering techniques to surface previously unidentified fraud cohorts.

Quantum-enhanced models are also integrated to provide enhanced capabilities, including sharper classification boundaries, richer multivariate pattern recognition, and expedited feature selection across extensive signal spaces. Knowledge graph and network analytics capabilities are then leveraged to transform an isolated flagged account into a mapped fraud network through multi-hop entity tracing.

03

Act & improve

Intelligence is only effective if it reaches the right operational teams without delay. Findings are presented via an analyst workbench designed for triage and transparency, rather than operating as an opaque system. This workbench can be augmented by an agentic AI assistant that enhances case data, responds to natural language queries, and formulates remediation guidance. This support facilitates rapid execution of remedial and preventative actions, such as fund, data, and property recovery, and fraud actor network disruption, upon case confirmation. Outcome labels from confirmed cases are fed directly to the fraud activity database, initiating automated model retraining. This iterative feedback loop ultimately enables autonomous agents to manage routine, low-complexity cases, thereby allowing human analysts to concentrate on more intricate challenges.

Built for how fraud teams actually work

- **Interdicts at any stage, not just the strike**. Early deployment of countermeasures during target selection or reconnaissance phases is more cost-effective and less disruptive than post-loss remediation. Waiting for the strike to respond is the most expensive option, not the safest one.
- **Calibrates response instead of reacting by reflex**. Not every outlier is fraud. The model is built to separate legitimate outliers from fraudulent ones, so good customers aren't punished for looking unusual while genuine attacks are more likely to be identified and addressed.
- **Sees across silos**. Cyber, transactions, operations, identity, and third-party signals are brought together into one unified view instead of five siloed ones, closing the gap fraud actors can exploit.
- **Moves at the adversary's pace**. As fraud actors adopt agentic AI and Fraud-as-a-Service to launch more attacks faster, the model is designed to adapt tactics and targets just as quickly, rather than waiting for the next scheduled model refresh.
- **Blends classical and quantum-enhanced models**. Proven classical ML handles the bulk of scoring and pattern-matching; quantum-enhanced clustering, classification, and optimization are layered specifically where classical approaches reach their limits, specifically in analyzing complex multivariate patterns, fuzzy classification boundaries, and feature selection across very large signal spaces.
- **Doubles as a business intelligence asset**. The same signal picture built to catch bad actors also surfaces the inverse: where an organization's good customers and suppliers actually sit and how they behave. This gives the business a clearer read on relationships worth investing in, not just ones to watch. Pointed forward, that same intelligence can be mapped against product and service roadmaps to flag where a new offering is most exploitable before it ships, instead of discovering the gap after fraud actors do.

Why this matters

Integrated defense

One connected view across domains instead of independent, partial ones.

Value-driven outcomes

Profit protection, revenue quality, key performance indicators (KPIs), and key risk indicators (KRIs) leadership can track across fraud prevention, detection, and operations.

Predictive posture

Anticipate emerging and resurging typologies before they cause losses, rather than discovering them afterward.

Compounding advantage

Institutional memory that carries lessons from one case, unit, or region to the next, so the program gets smarter over time instead of starting over with every new tactic.

Offensive action, not just detection

Disrupt fraud actor networks and pursue recovery of funds, data, and property once a scheme is confirmed.

Responsible by design

Operational velocity should not compromise trust. The model is engineered to maintain human oversight, facilitating a "human in the loop" over all recommended actions. This is supported by explainable outputs and well-defined escalation and change-control frameworks. An underlying governance layer is developed for this purpose, incorporating role-based access control, single sign-on, and multi-factor authentication protocols preceding any data movement. Furthermore, audit logging, data lineage tracking, and continuous bias monitoring are integrated from the beginning. Operational dashboards provide real-time tracking of KPIs, Return on Investment (ROI), and KRIs, featuring automated escalation mechanisms. The architectural design is also aligned with pertinent compliance frameworks and standards, such as SOC 2 and NIST.

Contact us



Shuba Balasubramanian

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



Mike Weil

Managing Director
Deloitte Financial Advisory Services LLP
miweil@deloitte.com

Deloitte.

This document contains general information only and Deloitte is not, by means of this document, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This document 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 document.

As used in this document, "Deloitte" means Deloitte Financial Advisory Services LLP, which provides forensic, anti-fraud, litigation, dispute, and computer forensic 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. Deloitte does not provide legal services and will not provide any legal advice or address any questions of law.

Copyright © 2026 Deloitte Development LLC. All rights reserved.