



GLOBAL FRAUD INTELLIGENCE REPORT 2026

AI, Identity and The Future of Fraud Defense



Contents

Executive Summary

The Threat Landscape

Part I: The Rising Scale And Cost Of Fraud

Chapter 1: Fraud's New Operating Model and The Identity Crisis	04
--	----

Part II: The Identity Crisis Behind Digital Deception

Chapter 2: Identity Fabrication: Synthetic Identities and Deepfakes	11
---	----

Chapter 3: Account Takeover 2.0	16
---------------------------------	----

Chapter 4: Mule Accounts	23
--------------------------	----

Part III: An Expanding Attack Surface Across Industries

Chapter 5: Financial Services (Banking, Fintech, NBFC)	27
--	----

Chapter 6: eCommerce, Quick Commerce, and Gig Platforms	33
---	----

The Economic and Regulatory Impact

Part IV: Regulation Is Raising The Stakes

Chapter 7: The True Cost of Fraud: Beyond the Loss Line	40
---	----

Chapter 8: The Global Compliance Landscape	45
--	----

Chapter 9: The Cost of Non-Compliance	49
---------------------------------------	----

Regional Risk Intelligence

Part V: The Global Reach Of Fraud Networks

Chapter 10: United States	53
---------------------------	----

Chapter 11: Europe and United Kingdom	57
---------------------------------------	----

Chapter 12: Southeast Asia	60
----------------------------	----

Chapter 13: Asia-Pacific	64
--------------------------	----

Chapter 14: Middle East and North Africa	69
--	----

The Fraud-Resilient Playbook

Part VI: AI Is Accelerating The Threat

Chapter 15: AI on Both Sides of the Battlefield	74
---	----

Chapter 16: Model Governance: Arms Race, Explainability and Regulatory Risk	77
---	----

Part VII: A New Playbook For Risk Leaders

Chapter 17: The Investment Thesis: From Cost Center to Competitive Advantage	81
--	----

Chapter 18: Ten Strategic Priorities for 2026-2027	84
--	----

Chapter 19: The Fraud-Resilient Enterprise Architecture Blueprint	87
---	----

APPENDIX

1. Bureau Unified Risk Decisioning Platform

2. Glossary

Executive Summary

Driven by organized rings and AI-powered tools, fraud has industrialized into a global US\$442 billion enterprise⁹. What was once a crime driven by individually motivated actors, where fraud could be contained by friction, managed through controls, and absorbed as an operational cost, has now evolved into coordinated attacks orchestrated by multiple actors across an interconnected fraud ecosystem. It is organized, technology-enabled, and built to scale. Adversaries now operate with product roadmaps and fraud-as-a-service models that outpace legacy institutional defenses.

This report analyzes current threat vectors, shifting regulations, and defense strategies to help decision-makers build resilient institutions.

Six Key Findings

The fraud ecosystem has industrialized.

Bureau's fraud intelligence network identified 14,000 organized fraud rings in H1 2026: 1 in 3 contained resurfacing identities, 1 in 4 spanned multiple industries, and the largest connected over 45,000 identities.

Identity has become the primary battleground.

Deepfakes, synthetic documents, and voice cloning have expanded attacks beyond onboarding into account access. FATF assessments reveal 90% of jurisdictions now flag identity fabrication as a major money-laundering risk.

Boards are measuring only part of the problem.

Direct losses are just the floor. The FTC estimates US consumer fraud losses reached up to \$195.9B (over 15x the \$12.5B reported in 2024). Factoring in abandoned journeys, brand impact, and regulatory liability, the true cost reaches 2-3x direct losses.

Regulatory liability is now a predictable business cost.

Compliance is an ongoing operating expense. UK PSPs reimbursed £243M to APP fraud victims in 15 months, while the EU AI Act carries fines up to €15M or 3% of global turnover. Strict frameworks (PSD3/PSR, SPF, UAE OTP ban, SAMA) phase in through 2027.

Yesterday's detection architecture cannot stop today's fraud:

Static, batch-processing models fail against real-time threats. In H1 2026, Bureau detected 21.7M+ account takeover attempts across 5.4B login sessions (~1 in 250 attempts).

Fraud resilience is becoming a competitive advantage and the window to build it is narrowing:

Modernized defenses drive growth. Institutions using Bureau's unified identity network are already seeing 30% faster onboarding, 83% faster sign-ups, and 65% fewer drop-offs.

As new mandates take effect through 2026-2027, advanced fraud capabilities will shift from a competitive edge to a baseline requirement. Because direct losses are only a fraction of total costs that include friction, brand damage, and compliance, institutions must build resilience now to maintain market advantage

SECTION 01

The Threat Landscape

CHAPTER 01

Fraud's New Operating Model and The Identity Crisis

Fraud's New Operating Model and The Identity Crisis

- **Fraud Has Industrialized:** Criminal networks now operate like industries, offering scalable dark-web products, SLAs, and AI-powered attack tools.
- **Three Forces Are Converging:** Plug-and-play fraud, autonomous AI agents, and real-time payments are lowering entry barriers and shrinking detection windows.
- **Massive, Undercounted Global Losses:** Fraud losses reached an estimated \$442 billion in 2025, with reported figures still understating the true cost.
- **The Failure of Point-in-Time Verification:** AI-generated forgeries and widespread PII leaks have made one-time credential checks increasingly ineffective.
- **A Shift to Continuous Identity Signals:** Effective prevention requires behavioral biometrics, cross-institutional identity graphs, and hardware-level cryptographic attestation.

14,000

organized fraud rings detected across Bureau's network

1 in 170

new accounts exhibited mule-like characteristics.

1 in 250

logins was an account takeover (ATO) attempt.

42%

of high-risk sessions were associated with devices linked to multiple user accounts.

The Industrialization Threshold

Three forces, as enumerated below, have converged to produce the current environment. Each would have been significant in isolation. But, together they have created a threat that is categorically different from what preceded it.

7.4 million

users exhibited suspicious multi-device behavior associated with potential promotional abuse.

5.8 million

ATO events involving suspected synthetic identities were flagged in just three months.

91 million+

high-risk sessions identified on Android 15 and 16.

The fraud economy developed a supply chain.

Instead of building capability, criminal networks rent it. Phishing kits, synthetic identity packages, mule networks, and deepfake generation are products that can be purchased; and are accompanied by documentation, update cycles, and service guarantees. The consequence is a decoupling of actor capability from attack capability: a low-sophistication operator with access to the right marketplace can now deploy an attack indistinguishable from the output of a skilled team.

As a result, for most attack types, the barrier to entry for sophisticated fraud has not just lowered, it has effectively disappeared. Bureau's fraud intelligence network identified roughly **one in every 170 onboarding applications as a suspected mule account** over the last five quarters. The geographic patterns remained consistent, showing a reusable fraud intelligence infrastructure rather than a response rebuilt after each disruption.

Generative AI moved from capability to operational deployment.

The FATF's December 2025 Horizon Scan has documented AI agents executing complete fraud campaigns without human intervention at each step; not as a future scenario, but as a present reality. The FBI has **logged AI-facilitated fraud as a standalone category for the first time** in 2025, and noted that even its figures were an undercount,

as most victims do not recognize AI involvement in the attacks targeting them. Therefore, now it is not a question of whether AI would transform the fraud threat, it is whether institutions can close the capability gap before the cost of not doing so becomes unrecoverable.


Real-time payment rails created a detection gap the fraud industry was built to exploit.

When funds move in seconds across UPI, FedNow, PayNow, FPS, and PromptPay, the intervention window collapses below the response time of any detection system designed for a previous era. This is an **architectural mismatch between infrastructure designed for speed and risk systems designed for accuracy**. APP fraud, pig-butcher scams, and real-time account takeover extraction are all built on the same principle: move faster than detection can respond. Suspected phishing activity on quick commerce platforms nearly doubled in two quarters,

per Bureau's fraud intelligence network; growth that outpaced eCommerce platforms over the same period. Bureau's survey of senior risk and fraud leaders across the Middle East banking, fintech, and adjacent sectors found faster, **instant payment adoption to be the most frequently cited driver of expected fraud growth into 2026**, ahead of AI-driven phishing, institutional data silos, and low consumer awareness. The pattern echoes what real-time rails have already produced elsewhere: detection compressed to the point of failure.

Phishing Activity Over Time

Suspected phishing events across industries and quarters



KEY TAKEAWAYS:

Quick Commerce continues to experience higher phishing rates compared to eCommerce across all quarters.

ECOMMERCE Phishing rates remain low in 2026 at 0.1% in both Q1 and Q2, down from 0.2% in Q4 2025.	QUICK COMMERCE Phishing rates increased in 2026, rising from 0.2% in Q4 2025 to 0.4% in Q2 2026.
--	---

SOURCE: BUREAU FRAUD INTELLIGENCE FRAUD NETWORK | DATA AS OF Q2 2026

The Scale of What This Industry Produces

Precise global loss figures are not available, as **underreporting is endemic**, classification is inconsistent across jurisdictions, and the most damaging fraud types often never appear in any loss report. What the data does establish is the direction and order of magnitude. Both are unambiguous.

For instance, INTERPOL estimated global fraud losses at \$442 billion in 2025⁴, projecting significant escalation as **AI continues to lower barriers to entry**. In the EU, Europol recorded \$64.1 billion in losses and named fraud the fastest-growing area of organized crime in the bloc. In the US, internet crime losses reached \$20.877 billion, up 26% in a single year.

Below are some reported figures from the most capable reporting systems in the world. They are undercounts.



Reported losses from the most capable reporting systems in the world. They are undercounts.

GEOGRAPHY	REPORTING BODY	PERIOD	REPORTED LOSS	YOY CHANGE
Global	INTERPOL	2025	\$442 billion (est.)	—
United States	FBI IC3	2025	\$20.877 billion	+26%
European Union	Europol	2025	\$64.1 billion	—
United Kingdom	UK Finance	H1 2025	£629.3 million	+3%
India	RBI	FY2025–26	₹48,021 crore (~\$5.6B)	+46.4%*
SEA (Singapore)	Singapore Police Force	2025	S\$913.1 million	—
Australia	National Anti-Scam Centre	2025	AUD \$2.18 billion	+7.8%
Middle East & North Africa	INTERPOL (member assessment)	2025–26	Not quantified	—

* INTERPOL's global estimate has no directly comparable prior-year baseline; this is the first edition to publish a consolidated global loss figure in this format. **Includes legacy cases reclassified under India's Supreme Court March 2023 directive. ***Singapore's decline followed coordinated mule restrictions under the October 2025 Facility Restriction Framework, a documented case of institutional intervention measurably outpacing the threat. All figures are directional indicators. True losses exceed what any reporting system captures.

The Identity Crisis

Every fraud attack starts with a claim about who someone is. The fraud industry's competitive advantage lies in making that claim convincingly; and the tools now available to do so have fundamentally compromised the verification model most institutions still rely on.

The FATF's February 2026 Cyber-Enabled Fraud paper identified AI-generated document forgery as a primary fabrication technique at operational scale.

The implications reach further than document fraud. The model of point-in-time verification, presenting a credential at a single moment to establish identity, was designed for a world where fabrication was technically difficult and expensive. That constraint no longer exists.



The Verification Trust Stack – 2026

The signals institutions rely on to verify identity are not equally defensible. The hierarchy has shifted. What once anchored identity assurance now sits at the bottom of the stack; what was once supplementary has become the most reliable layer available.

Ranked most to least reliable as a standalone identity signal:

Cryptographic device attestation Verification anchored to hardware, not presented credentials.
Cross-institutional identity graph matching Effective against synthetic identities when consortium data is available.
Behavioral biometrics and continuous session signals Interaction patterns that cannot be replicated by document presentation.
Device intelligence and network context Hardware environment, OS integrity, network telemetry.
Liveness detection with active challenge-response Reliable only when cryptographically anchored and session-unique..
Document verification Compromised as a standalone signal at scale.
Knowledge-based authentication Effectively deprecated; the PII it requires is available on dark-web markets.

The response, thus needed, is a shift from point-in-time verification to continuous identity assurance: ongoing behavioral monitoring across the customer lifecycle that does not treat onboarding as a settled question. This is a structural response to an adversary ecosystem that has industrialized the fabrication of the credentials point-in-time systems are designed to check.

Institutions still anchoring identity assurance in document verification are defending a position the fraud industry has already outflanked. But the next layer of defense carries its own blind spots. Device intelligence is only as good as how it's read, and **device age alone is not a reliable signal**.

The Stakes

Across the US, UK and Europe, India, Southeast Asia, and Australia, different reporting systems, methodologies, and classification frameworks point to the same conclusion⁸. The Middle East and North Africa presents a distinct posture: **fraud risk currently rated moderate relative to the global average**, but with scam-compound infrastructure actively migrating into the region and a fintech adoption curve that is creating exposure faster than institutional defenses are being built.

Across all regions, fraud is no longer a line-item risk managed within existing operational budgets. It is a strategic threat to institutional viability, customer trust, and regulatory standing, prosecuted by **an adversary ecosystem that is better organized, better resourced, and better adapted** to the current environment than at any point in the history of financial crime.



BUREAU CLOSE

Each point solution buys time. **The adversary ecosystem is designed to route around point solutions and exploit the gaps between them.** The question to answer, therefore, is not whether your institution is exposed; it is where the exposure is greatest, how it will evolve, and what an adequate response architecture looks like.

SECTION 01

The Threat Landscape

CHAPTER 02

Identity Fabrication:
Synthetic Identities
and Deepfakes

CHAPTER 03

Account Takeover 2.0

CHAPTER 04

Mule Accounts

Identity Fabrication: Synthetic Identities and Deepfakes

- **Generative AI Reduces Identity Fabrication Costs:** Criminal networks create convincing identities at near-zero cost, weakening traditional verification.
- **Synthetic Identities Lack a Victim Signal:** Fabricated profiles build credible histories before coordinated bust-outs, leaving no victim alerts or dispute triggers.
- **AI Defeats Conventional Media Verification:** Deepfakes, voice clones, and forged documents increasingly bypass static identity and liveness checks.
- **Data Fragmentation Masks Multi-Lender Fraud:** Siloed institutions cannot see cross-lender activity, allowing synthetic identities to borrow undetected.
- **Defense Requires Continuous Multi-Signal Assurance:** Protection must move beyond point-in-time checks to continuous monitoring, behavioral biometrics, identity graphs, and hardware attestation.



How Synthetic Identities Are Built

The key distinction between a synthetic identity and a stolen identity is the presence of a real victim. Stolen identities have victims who notice unauthorized activity and trigger investigations. Synthetic identities have no victims because the person they represent does not exist.

Constructing a synthetic identity follows a consistent pattern. Fraudsters combine genuine identity fragments: PII sourced from dark-web data markets, often belonging to individuals unlikely to monitor their credit, with fabricated elements such as a name, address, employment history, and a digital footprint designed to pre-date the fraud operation. Generative AI now automates much of this process, creating synthetic profiles that are internally consistent, biographically coherent, and visually credible.

The constructed identity enters the financial system through its lowest-friction access point, say a secured card, a retail account,

The **constructed identity enters the financial system through its lowest-friction access point**, say a secured card, a retail account, or a buy-now-pay-later facility. Initial behavior is unremarkable by design: on-time payments, appropriate utilization, and no anomalies. However, over a few months and in some cases years, the identity accumulates credit history, builds relationships across multiple institutions, and increases available credit. Generative AI has made this process faster and more adaptive. Synthetic identities can now learn from failed attempts, vary their profiles to avoid duplication, and target institutions where detection controls are the weakest.

GenAI is now identified as a significant accelerant of this threat²³, making synthetic identities harder to detect, easier to iterate, and more capable of scaling across the financial system.

The final stage is the bust-out: every available credit line across every linked institution is drawn down simultaneously, & then the identity disappears. No victim files a complaint. No dispute triggers investigation.

The losses are absorbed, often as credit risk, and in many cases never classified as fraud at all.

Synthetic identities are also increasingly the entry point for account takeover, not just credit fraud. Bureau's fraud intelligence network recorded **account takeover events linked to suspected synthetic identities nearly tripling in a single quarter**: evidence that fabricated identities are being built for reuse across attack types, not deployed once and discarded.

Why Detection Failure Is Structural, Not Accidental

Delayed detection is a defining characteristic of synthetic identity fraud, rather than a result of institutional oversight. It does not get detected quickly because the **fraud is architecturally designed to look like legitimate customer behavior** for its entire useful life. Three structural conditions make early detection genuinely difficult, and none can be addressed through a faster or more sophisticated point-in-time checking.

Conventional fraud detection relies heavily on customer-reported anomalies, such as unauthorized charges, mismatched credit entries, or complaints that open investigations. Synthetic identities generate none of this. The absence of a living person monitoring the identity is not a gap in the attack. It is a core feature of the design.


This risk extends even to platforms built around trust between users. Bureau's fraud intelligence network **flagged more than 20,000 suspected synthetic or scam accounts on matrimony platforms in H1 2026**; identities created specifically to exploit relationship-based trust rather than financial access.

During the cultivation phase, synthetic identities pay on time, maintain appropriate balances, and earn real credit scores. They are not just simulating good credit behavior, but also exhibiting it.

As a result, no rules-based system flags behavior that is indistinguishable from a legitimate thin-file customer trying to build credit history.


The same synthetic identity can borrow from multiple lenders at the same time, but this pattern is only visible when data is shared across institutions in real time. Each lender sees only part of the picture. The full fraud pattern only emerges when data from multiple institutions is combined.

The FATF's February 2026 Cyber-Enabled Fraud paper²⁴ identifies institutional data fragmentation as a primary structural vulnerability, noting that AI-generated identities can sustain multi-institution relationships over extended periods before exploitation, precisely because no single institution has visibility of the full picture.



Deepfake Video Attacks, Voice Cloning, and AI-Generated Document Forgery

Synthetic identities once failed at the verification stage. A mismatched photograph or an inconsistent document could expose the fraud. However, **generative AI has now closed that gap**, enabling synthetic identities to appear consistent across every layer of the conventional identity verification process.



AI-generated document forgery

Has moved from a detectable risk to an operational reality, one that FATF now treats as an active threat. It strikes where institutions trust most: customer identification, defeated before any other check runs.

Deepfake video attacks

Defeat the liveness verification layer. Fraudsters inject pre-recorded or AI-generated video into live verification sessions, presenting a synthetic face to the verification system without triggering detection. The FATF's December 2025 Horizon Scan²⁵ named deepfakes deployed at scale for KYC bypass as a present operational threat. It also confirmed that the tools enabling these attacks are accessible without high technical skill, extending their reach across the full spectrum of fraud actors, not just sophisticated ones.

Real-time voice cloning

Extends fabrication to the audio channel. Europol's EU-SOCTA 2025²⁶ identified **voice cloning and live video deepfakes as tools actively deployed in CEO fraud and romance scams**. AI-powered voice synthesis enables fraudsters to impersonate executives, colleagues, and authority figures in real-time over phone and video channels. Practitioners on the ground report the same pattern. Among senior risk and fraud leaders across the Middle East polled by Bureau, the majority observed at least occasional real-time impersonation, device spoofing, or emulator-based attacks against their customers, with a meaningful share describing it as very frequent.

The Bank of England's Financial Stability in Focus paper²⁷ on AI frames the compound risk directly.

It says the capability of generative AI models to produce deepfakes, alongside highly personalized text, increases the ability of those intent on fraud to manipulate both employees and retail customers, a risk that extends beyond the identity verification layer into internal authorization processes and customer-facing channels.

The result is a verification environment in which every conventional layer, from document, to photograph, video, and voice, is subject to fabrication at a quality level that static, point-in-time checking cannot reliably defeat. The adversary does not need to beat every control. It only needs to find the weakest combination.

The Defense Path: From Document Checks to Continuous Identity Assurance

The defense challenge is clear. If every static identity element can be fabricated at scale, **verification systems that rely on checking those elements become inherently vulnerable**. Better verification tools address only part of the problem. The real challenge is architectural.

Two shifts, as described below, are required.

From single-layer to multi-signal verification

The signals hardest to fabricate simultaneously are behavioral and device-level, not document-level. A convincing deepfake does not replicate the cryptographic attestation of specific hardware. A synthetic identity with a perfect credit history does not naturally produce the behavioral session patterns of a real person interacting with a familiar interface. The defense architecture must treat these signals as primary, and document verification as one input among several, not the anchor it has historically been.

The signals institutions should be building toward, ranked by resistance to fabrication at scale, include:

Cryptographic device attestation Verification anchored to hardware through proof, not presented credentials.	Cross-institutional identity graph matching The only signal that makes the bust-out pattern visible; requires consortium data infrastructure.	Behavioral biometrics and continuous session signals Interaction patterns built over time, not replicable at a single verification checkpoint.
Device intelligence and network context Hardware environment, OS integrity, network telemetry.	Liveness detection with active challenge-response Reliable only when cryptographically anchored and session-unique.	Document verification Useful as one layer in a multi-signal architecture; insufficient as a primary control.

From point-in-time to continuous identity assurance


Onboarding verification answers one question: does this person appear to be who they claim to be, with these credentials, at this moment? It does not answer the more consequential question: does the behavior of this account, over time, remain consistent with a legitimate customer? Continuous identity assurance, ongoing behavioral monitoring across the full customer lifecycle, is what can detect synthetic identities during the cultivation phase, before the bust-out. That requires rearchitecting how identity is treated across the full customer relationship, not just at the front door.

BUREAU CLOSE

Synthetic identity fraud and AI-enabled verification bypass are the predictable output of adversaries who studied those systems and engineered around them. The response cannot be incremental: it must **rethink what verification is designed to establish, and whether a check performed once at onboarding is sufficient** in an environment where the fabrication of credentials has become industrialized.

Account Takeover 2.0

- **ATO Exploits Post-Authentication Sessions:** Attackers bypass MFA and passwords through stolen sessions, infostealers, and social engineering, defeating login-time checks.
- **Modern ATO Follows a 4-Stage Chain:** Acquisition, reconnaissance, access pivoting, and rapid extraction occur before traditional alerts trigger.
- **Primary Attack Vectors Target Session and Device Layers:** Infostealers, SIM swaps, and MFA fatigue target sessions, devices, and user behavior, not credentials.
- **Continuous Behavioral Monitoring Is Essential:** Real-time session telemetry reveals anomalies that point-in-time login checks miss.
- **Channel-Specific Signals Reveal Hijacked Sessions:** Device emulators, app tampering, fingerprint mismatches, and cloned voices expose hijacked sessions.



The Modern ATO Kill Chain

The architecture of a modern account takeover operation follows a consistent logic; and that logic has shifted decisively toward post-credential compromise. Bureau's fraud intelligence network found that **more than one in every eight high-risk account takeover sessions carried indicators of social engineering**, underscoring the growing role of human manipulation alongside technical compromise.

ACQUISITION

Credentials, session cookies, and authentication tokens are acquired through one of three primary routes:

- phishing that harvests active session tokens rather than just passwords;
- infostealer malware that silently extracts browser cookies from infected devices; or
- social engineering that manipulates the legitimate account holder into surrendering access directly.

RECONNAISSANCE

Once access is established, sophisticated actors do not immediately extract value. They observe: mapping account balances, linked accounts, payment limits, beneficiary lists, and authentication configurations. **This phase is largely invisible to detection systems calibrated to flag anomalies at login** rather than monitoring for low-velocity behavioral deviation post-authentication.

PIVOT

Contact details are changed. New beneficiaries are added. Authentication recovery options are redirected. The **attacker establishes durable control** that survives a password reset by the legitimate account holder.

EXTRACTION


Funds are transferred, stored payment credentials are exploited, loyalty points are liquidated, or the compromised account is resold to downstream fraud actors. **The extraction window is measured in hours**, often completed before any detection system has generated a flag.

Since January 2025, the FBI received more than 5,100 complaints specifically reporting ATO fraud via impersonation of financial institution support²⁹, with losses exceeding \$262 million, a figure covering one attack modality within a broader ATO landscape that extends across all channels and attack types.

The scale of account takeover risk is significant. Bureau's fraud intelligence network identified **more than 21.7 million account takeover attempts across 5.4 billion login sessions** in H1 2026, roughly **one in every 250 login attempts**. Risk levels also increased through the quarter, rising from 0.29% of sessions in April to 0.49% in June.

Account Takeover (ATO) Attempt Rate by Month - 2026

ATO ATTEMPTS AS A PERCENTAGE OF TOTAL EVENTS:



Why ATO Is Now a Behavioral Problem, Not a Credential Problem


The defining characteristic of modern account takeover is that the most damaging attacks bypass the authentication layer entirely. A stolen session cookie does not require a username, password, or MFA code. It is an authenticated state; and to the system it targets, **it is indistinguishable from the legitimate user who created it.**

This is an architectural consequence of how modern authentication works. When a user successfully authenticates, the system issues a token confirming that authentication has occurred. The token is what governs subsequent access.

The infostealer ecosystem has **industrialized the theft of those tokens at scale**, effectively allowing attackers to step into authenticated sessions without ever touching the credential that created them.

The implication for defense is structural.

A fraud detection model built around login-time signals: device fingerprint, IP reputation, credential validity, does not see this attack at all. The session is clean. The credentials were never compromised. The MFA challenge was never triggered. The attack exists entirely in the space between successful authentication and the end of the session.



SIM Swap, MFA Fatigue, and Session Hijacking: The Primary Vectors

ATO usually arrives through three distinct access routes, each targeting a different weakness in the authentication and session management chain.

Session hijacking via infostealer

Infostealer malware functions as an industrial infrastructure for session token theft. It distributes stolen cookies at scale, through dark-web marketplaces that operate with the efficiency and resilience of legitimate software platforms. Session cookies harvested from an infected device can be imported into an attacker-controlled browser, **bypassing all login controls including MFA**, because the system interprets the imported session as a continuation of the legitimate user's authenticated state.

Device reuse is a key indicator of ATO risk. Bureau's fraud intelligence network found that **more than 42% of high-risk ATO sessions involved devices linked to multiple accounts**, and over one-third showed signs of app tampering, cloned applications, or unofficial installs. The pattern shows that **attacks are increasingly targeting the device and session layer**, not just credentials.

SIM swap



By porting a victim's mobile number to a fraudster-controlled SIM card, attackers gain access to every SMS-based one-time password and account recovery code the victim's financial institution sends. **SIM swap remains a primary enabler of high-value ATO**, particularly in markets where SMS-based MFA remains the dominant second factor for financial services authentication. Europol's IOCTA 2026 identifies telecoms infrastructure compromise as a persistent entry point for subsequent fraud chain exploitation.

MFA fatigue



Rather than defeating MFA technically, fatigue attacks exploit the human layer, sending repeated push authentication requests to a legitimate user **until they approve one in error**, believing it is a system glitch. The FATF has identified social engineering that manipulates legitimate account holders into surrendering access as a documented operational vector, one that bypasses technical controls by targeting human behavior rather than system vulnerabilities.

Post-Authentication Behavioral Detection: The Frontier

The response to the behavioral ATO problem requires a behavioral answer. If the attack is invisible at authentication, **detection has to happen after authentication**, in the session itself, across the customer lifecycle, calibrated to a baseline of what the legitimate account holder actually does. Post-authentication behavioral detection operates on the premise that even when an attacker successfully inherits a session, their behavior within that session will deviate from the established pattern of the legitimate account holder. The deviation may be subtle: a different navigation path, an unusual sequence of actions, or interaction patterns inconsistent with the device and time of day. But it is detectable if the right baseline exists and the right signals are monitored.

The signals that matter for post-authentication detection differ from those used at login. They include:

- Session-level behavioral patterns: how a user navigates, types, and interacts with interface elements
- Consistency between the current session's device telemetry and the account holder's established device profile
- Action sequences that deviate from the account's historical transaction pattern; and
- Velocity anomalies in account modification activity: changes to contact details, beneficiary additions, authentication setting updates, inconsistent with legitimate account management behavior.

Device Signal Patterns Across Channels and OS

The device and session provide important signals that go beyond whether the credentials are valid. Indicators that warrant closer scrutiny include:

Across mobile channels:

a mismatch between the device OS and the established device profile for the account;
a recently changed SIM or IMEI inconsistent with the account holder's device history; and
session initiation from a device emulator or virtual environment rather than a physical device.

Across web channels:

browser fingerprint inconsistencies between the current session and the account's established access pattern;
Presence of automation signatures: request timing, scroll behavior, interaction cadence, inconsistent with human navigation; and
network context anomalies, including residential proxy infrastructure and IP geolocation inconsistencies.

Bureau's fraud intelligence network recorded more than 91 million high-risk sessions on Android 15 and 16 devices alone in Q2 2026, evidence that **account takeover increasingly targets the devices most widely used by legitimate customers**, not outdated or unsupported hardware. iOS sessions, by comparison, carried high-risk rates of roughly 0.2% across most versions.

Across IVR and telephony channels, voice authentication bypass via cloned audio and call routing through VoIP infrastructure, masking the true originating number,

are established typologies documented by Europol³² in the context of AI-enabled social engineering deployed against customer-facing financial service channels.

The same principle applies across every channel: **authentication confirms entry, but it does not establish trust for the rest of the session.** That requires continuous monitoring of device, behavioral, and network signals after authentication. This is the control layer today's threat landscape demands, yet one that many institutions have yet to implement.


BUREAU CLOSE

Account takeover attack begins after authentication has succeeded. The **threat now exists within the session**, in the period between a successful login and the end of the interaction. Institutions that focus ATO defenses on the login checkpoint are protecting only the point of entry. **Effective detection now depends on continuous monitoring** of device, behavioral, and network signals **throughout the session.**

ATO Device (Mobile): High Risk Sessions by Budget Segment

Risk is broadly similar across known device segments (~1% high-risk rate), suggesting device affordability is not a strong fraud indicator.

% RISKY SESSIONS




High Risk Session Rate by Browser

Mainstream browsers and in-app browsing environments (>100,000 sessions)

BROWSER

% OF RISKY SESSIONS (RISK SCORE 70)



In-app browsing environments recorded materially higher fraud rates than traditional browsers. Safari (0.48%), Facebook App (0.41%) , and Instagram App (0.31%) exhibited the highest risk rates among high-volume browsers, while Chrome — despite the largest share of traffic at 112.8 million sessions — maintained a low rate of 0.1%.

Mule Accounts

- **Mule Networks Enable Fraud Economics:** Fraudsters use mule accounts to cash out and move stolen funds before institutions can intervene.
- **Deception Drives Mule Recruitment:** Job scams recruit mules whose onboarding and early activity resemble legitimate customers.
- **Reporting Gaps Empower Mule Networks:** Without cross-industry data sharing, networks can reuse recruits across institutions.
- **Mule Detection Demands Lifecycle Signals:** Detection requires ongoing behavioral, transactional, and network monitoring.
- **Network Interruption Requires Cross-Institutional Action:** Disruption requires shared intelligence, multi-bank analysis, and coordinated enforcement.

The Mule Account's Role in the Fraud Economy

Mule accounts occupy a specific and essential position in the fraud supply chain. Between the point of theft, such as a social engineering attack, an account takeover, or an APP fraud transfer, and the point of usable criminal proceeds lies a layering process designed to put distance between the crime and the criminal. Mule accounts are that distance.

Funds enter a mule account from a victim. They are **rapidly moved onward to another mule account, to a cryptocurrency exchange, or to an overseas transfer**, thereby fragmenting the trail and compressing the intervention window.

In many cases, the entire layering sequence completes within minutes of the original fraud, exploiting the same real-time rail speed that enables the fraud itself.

INTERPOL's Operation First Light 2026³³, which ran across 97 countries from January to April 2026, targeted precisely this infrastructure: the money laundering networks that process the proceeds of social engineering fraud. The operation resulted in 5,811 arrests, \$293 million in illicit assets intercepted, and 31,014 bank accounts blocked, a clear evidence of the scale and geographic distribution of mule account operations globally.

INTERPOL · OPERATION FIRST LIGHT 2026³³

97

countries, January to April 2026

5,811

arrests

\$293M

illicit assets intercepted

31,014

bank accounts blocked

How mules are recruited

Recruitment runs the full spectrum from coercion to willing participation, but the largest volume comes through deception, including fake job offers, fabricated business opportunities, and online marketplace overpayment schemes that frame account forwarding as legitimate work. Cifas Fraudscape 2026 identifies social media as the primary recruitment channel, and notes that a fifth of recruited individuals do not believe what they are doing is illegal. That framing matters for detection: the majority of active mule accounts at any given time belong to people who believe they are employed, which means their onboarding behavior, account usage patterns, and responses to bank queries mirror those of legitimate customers.

Mule-as-a-Service: Organized Networks at Scale

At the organized end of the spectrum, mule account management has become a service. Criminal networks maintain **inventories of verified mule accounts across multiple institutions**, cycling through them in response to closures, flagging accounts for specific transaction volumes and corridors, and replacing burned accounts with fresh recruits. The operational logic mirrors the fraud-as-a-service model: specialization, redundancy, and rapid response to disruption.

The scale of the institutional response required to contain this infrastructure is significant. According to the FCA's December 2025 findings³⁵, 25 UK firms offboarded 194,084 money mule accounts between January 2022 and September 2023, an average of more than 9,700 each month.

Yet only 37% of those cases were reported to the National Fraud Database. The remaining 63% reporting gap gave organized mule networks room to operate undetected. In nearly one-third of these remaining cases, the FCA found that firms should have reported them. As a result, most offboarded mule accounts were not flagged across the industry, allowing the same individuals to open accounts at other institutions and continue operating.

FCA · DECEMBER 2025 FINDINGS

"25 UK firms offboarded 194,084 money mules over 20 months, but 63% were never reported to the cross-industry database. The same individuals could immediately open accounts elsewhere."

OFFBOARDED

194,084

PER MONTH

9,700+

REPORTED TO NFD

37%

REPORTING GAP

63%

Detection Signals: Behavioral, Transactional, and Network

Mule detection is harder than fraud detection for a specific structural reason: at the point of account opening, and often for the initial period of activity, mule accounts are indistinguishable from legitimate customer accounts. The fraud has not yet happened. The mule has not yet been activated. There is nothing to detect.

Detection becomes possible when patterns emerge across three signal categories, all of which require ongoing monitoring rather than point-in-time checking.

BEHAVIORAL SIGNALS

They reflect the mule account's actual function: receiving and rapidly forwarding funds, rather than the transaction patterns of a genuine customer. Indicators include:

- Incoming transfers from unknown or unrelated parties inconsistent with the account holder's stated purpose
- Rapid onward transfer of received funds, often within minutes;
- Transaction amounts and frequencies inconsistent with the account holder's income or account type;
- Activity concentrated in narrow time windows, reflecting coordinated network behavior.


They reflect the structure of the layering process itself. reflect the structure of the layering process. Accounts that receive round-number transfers, immediately forward the majority of received funds to accounts at different institutions or in different jurisdictions, and maintain near-zero balances between activation events follow a pattern that distinguishes mule behavior from legitimate high-velocity transaction activity.

Together these distinguish mule behavior from legitimate high- velocity transaction activity.

NETWORK SIGNALS

These are the most diagnostic and the least accessible to any single institution. The mule account does not operate in isolation. It receives from one node in the network and forwards to another. Mapping those flows across institutions **reveals the network structure**, identifying accounts that function as mule coordinators, routing nodes, or final extraction points. This analysis is only possible with cross-institutional data access. Bureau's fraud intelligence network shows the scale and reach of organized fraud structures: **one in four identified fraud networks extended across multiple industries** rather than remaining within a single platform or sector. Mule accounts operate within these networks as reusable nodes, not isolated incidents to be closed individually.

Quarterly Mule Account Detection Rate



SINGAPORE · FACILITY RESTRICTION FRAMEWORK³⁶, OCT 2025

Rather than closing individual accounts, it placed 550 money mules, 801 SIM card mules, and 51 corporate entities under simultaneous banking and telecom restrictions, disrupting the network rather than trimming its edges. Scam losses fell 19% in 2025.

INDIA · MULEHUNTER.AI

India's RBI has built a cross-institutional infrastructure with MuleHunter.AI³⁷, which is now running across 23 banks, a detection rate no single institution could achieve through manual review alone.

Why Mule Detection Requires a Different Model

FRAUD DETECTION ASKS

Is this transaction legitimate?

Optimized for transaction-level anomaly identification: a single event that deviates from expected behavior. A model calibrated for fraud generates high false-positive rates on mule accounts, because individual mule transactions are structurally unremarkable.

MULE DETECTION ASKS

Is this account being used as infrastructure for someone else's fraud?

Requires account-lifecycle analysis: patterns visible only across time, across multiple transactions, and across the relationships between accounts at different institutions.

The FCA's December 2025 review found that every firm recognized the National Fraud Database as a critical tool for detecting money mules. Yet only 37% of confirmed mule accounts were reported. The gap reflects the operational burden of reporting, the absence of mandatory reporting requirements, and the view that closing an account is enough. It is not.

Disruption and Response

Detecting a mule account is only the first step. The response determines whether the detection has any impact on the network it has identified. Three response mechanisms operate at different levels of the problem.

Account restriction

Removes one node from the network. Necessary but insufficient in isolation, since organized mule networks replace closed accounts faster than institutions can close them.

MECHANISM 02

Cross-industry reporting

Extends a single institution's detection to the full consortium. When a confirmed mule is reported to a shared database, every member can act on that intelligence before the same individual opens an account elsewhere.

Coordinated law enforcement action

Operating on network intelligence rather than individual accounts, is the mechanism that matches the scale of the problem.

Using INTERPOL's I-GRIP (Global Rapid Intervention of Payments) mechanism³⁸, authorities in Singapore and Oman stopped a \$6.6 million business email compromise (BEC) transfer before the funds cleared. The operation showed that real-time, cross-border payment intervention is possible when the supporting institutional and regulatory infrastructure is in place.

BUREAU CLOSE

Mule accounts are the mechanism that **makes fraud economically viable at scale**. An institution that detects fraud but does not detect and disrupt the mule infrastructure processing its proceeds is solving half the problem, and **possibly even funding the next attack with the half it leaves running**. Effective mule disruption requires cross-institutional data sharing, AI-driven lifecycle analysis, and coordinated law enforcement engagement. None of those three things works in isolation.



SECTION 01

The Threat Landscape

CHAPTER 05

Financial Services: Banking,
Fintech, NBFC

CHAPTER 06

eCommerce, Quick Commerce,
and Gig Platforms

Financial Services (Banking, Fintech, NBFC)

- **Real-Time Rails Compress Intervention Windows:**
Instant payments eliminate detection timeframes, forcing institutions to evaluate risk before authorizing transfers.
- **Low-Friction Onboarding Accelerates KYC Bypass:**
Speed-optimized sign-ups allow GenAI synthetic documents and deepfakes to infiltrate digital lenders and BNPL platforms.
- **Account Takeovers Hijack Authenticated Sessions:**
Impersonation, SIM swaps, and session theft bypass login checks, giving fraudsters direct control to extract funds.
- **Synthetic Identities Fuel Credit Bust-Outs:** Fabricated profiles cultivate credit lines across multiple lenders before executing coordinated, high-value drawdowns.
- **BEC Targets Corporate Payment Authorizations:**
Business Email Compromise drives massive corporate losses by manipulating human authorization processes rather than technical systems.

Business Email Compromise: Corporate Banking's Most Expensive Problem

BEC remains the highest dollar-loss fraud category targeting financial institutions and their corporate clients. In 2025, the FBI IC3 recorded 24,768 BEC complaints with total losses of \$3.046 billion . This was the second-largest loss category in the report, behind only investment fraud, and an increase from \$2.77 billion in 2024.

\$3+ B

BEC losses, 2025 — the second-largest loss category in FBI reporting

24,768

BEC complaints that FBI IC3 recorded

\$2.77B

BEC losses, 2025 from 2024

Instead of technical compromise, BEC relies on the manipulation of authorization processes, such as **impersonating an executive, a vendor, or a legal representative to redirect payments, change beneficiary details, or extract sensitive financial data**. The attack surface is human and process-level, which means technical fraud controls provide limited protection and the intervention window is concentrated in the payment authorization moment itself. The risk to financial institutions is amplified by their dual role. They are both the channel through which BEC payments are executed and the institutions that bear the operational and reputational costs when those payments are disputed.

Synthetic Identity Bust-Out in Consumer Lending and BNPL

Synthetic identity fraud poses a particular risk to unsecured consumer lending. Within financial services, its impact is concentrated in two products: unsecured personal credit and buy-now-pay-later (BNPL) facilities.

Both share the same risk architecture: minimal physical interaction, rapid approval decisions, and a credit assessment model designed for speed rather than depth of identity verification. Synthetic identities exploit that architecture in the same way across both products. **First cultivating a thin credit history over time and then maximizing available credit across multiple lenders in a coordinated bust-out.**

The scale of the BNPL market has amplified this exposure. According to an estimate⁴⁰, total BNPL transaction value reached \$70 billion in 2025, representing approximately 20% annual growth since 2021.

The regulatory landscape further shapes the risk. The CFPB withdrew its BNPL interpretive rule in 2025, leaving BNPL providers without the consumer protection and reporting framework that applies to traditional credit products. Fraudsters have been quick to exploit this difference.

Digital Lending Fraud and KYC Bypass at Neobanks

Neobanks and digital lending platforms built their market position on frictionless onboarding: the same architecture that synthetic identity and deepfake-enabled KYC bypass exploit most effectively. The FATF identifies the low-friction digital onboarding model as a primary structural vulnerability across the financial sector globally,

with AI-generated documentation defeating conventional customer identification procedures at operational scale.

approximately 20% annual growth since 2021

Bureau's fraud intelligence network shows that mule account activity remains elevated. **Suspected mule accounts among banking onboardings nearly doubled between Q2 and Q4 2025**, before stabilizing in H1 2026 at a rate still roughly one-third higher than at the start of the period. The sustained level indicates that mule recruitment continues despite stronger detection, highlighting the **need for continuous, lifecycle-level monitoring beyond onboarding KYC.**

In India, the RBI's supervisory response reflects the scale of the problem. In early 2025, more than 2,000 unauthorized digital lending apps were flagged for removal from major app stores. These platforms operated outside the regulatory framework while offering credit to consumers. Their identity verification and credit assessments did not meet the standards required of licensed lenders. The RBI's Annual Report 2025-26⁴² signals a structural shift in supervisory philosophy: having AML and KYC policies is no longer sufficient, institutions must now demonstrate through evidence that their controls work in practice.

Senior fraud and risk leaders across the Middle East banking and fintech, polled separately by Bureau, identified **card-not-present fraud, remote banking and PSP fraud, money mules, and device spoofing as the vectors that grew fastest in 2025**, with no single category dominating. This suggests institutions in the region are being tested across the full attack surface rather than one point of failure.

The implication for global neobanks is clear. A KYC framework designed for legitimate thin-file customers will also accept well-constructed synthetic identities. Distinguishing between the two requires multiple risk signals and continuous behavioral monitoring, not document checks at onboarding alone.

Account Takeover in Banking Channels

Account takeover (ATO) poses a particularly high risk in banking because of the concentration of financial value in customer accounts. Once an attacker compromises an authenticated session, post-authentication authorization processes can be exploited to initiate high-value transfers that are often difficult to recover.

As reported earlier, the FBI received more than 5,100 account takeover (ATO) complaints involving fraudsters impersonating financial institution support staff, with reported losses exceeding \$262 million since January 2025. In these attacks, fraudsters pose as bank employees through phone calls, text messages, or fraudulent websites to convince customers to share authentication credentials or approve fraudulent transactions.

In the UK, account takeover cases rose 6% in 2025 to over 78,000⁴³, driven primarily by a 38% increase in telecoms-sector filings as SIM swap activity escalated. Mobile phone accounts, online retail accounts linked to financial products, and personal credit cards collectively account for 90% of all account takeover filings.

UK ATO cases in 2025, up 6%, driven primarily by a 38% increase in telecoms-sector filings as SIM swap activity escalated

of all UK ATO filings sit in mobile phone accounts, online retail accounts linked to financial products, and personal credit cards

\$262M

reported losses across 5,100+ FBI complaints where fraudsters impersonated financial institution support staff




UPI and Real-Time Payment Fraud in Fintech Rails

India's Unified Payments Interface has scaled to become the world's most used real-time payment system by transaction volume. This scale also makes it both a landmark financial inclusion achievement as well as one of the most targeted payment rails by fraud actors. The RBI's Annual Report 2025-26⁴⁴ points to a shift in the fraud landscape.

Digital payment fraud fell to 293 reported cases in FY2025-26, while advances fraud rose sharply to ₹40,774 crore. The trend suggests that as UPI fraud controls have matured, **sophisticated fraudsters are increasingly targeting higher-value lending fraud instead of high-volume payment fraud.**

293 reported cases



The pattern is not unique to India. Across real-time payment systems globally, such as the UK's Faster Payments Service, the US FedNow network, Singapore's PayNow, or Thailand's PromptPay, the fraud typology follows the same logic: **exploit the compression of the intervention window that real-time settlement creates.** The most successful attacks are completed before detection systems can generate an alert. Effective defense, therefore, requires risk assessment before a payment is authorized, not after it has been processed.

NBFC-Specific Application Fraud

Non-banking financial companies (NBFCs) occupy a distinct position in the fraud landscape. They often serve customers underserved by traditional banks, rely on lighter verification processes, and, in many markets, operate under less stringent regulatory requirements than licensed banks. This combination makes them **a preferred target for application fraud**, particularly synthetic identity bust-out and income misrepresentation at scale.

In India, the RBI's supervisory actions in FY2025-26 reflect the sector's exposure directly: enforcement actions for KYC non-compliance, digital lending guideline violations, and inadequate fraud controls were among the most common categories of regulatory action against NBFCs in the period. The RBI's stated agenda for FY2026 includes detailed thematic reviews on IT governance and cyber risk across all supervised entities, with NBFCs explicitly included in scope⁴⁵.

The risk for NBFCs extends beyond direct fraud losses. As markets introduce mandatory fraud compensation frameworks, such as the UK's APP reimbursement model and India's proposed unauthorized transaction compensation framework, NBFCs and fintech lenders that process payments or hold customer funds are likely to face similar liability obligations as banks. Many, however, have yet to invest in comparable fraud detection capabilities.

BUREAU CLOSE

Financial services is the sector where the fraud economy's full architecture is most visible; and most consequential. BEC exploits corporate authorization processes. Synthetic identities exploit credit underwriting models. KYC bypass exploits digital onboarding. ATO exploits authentication infrastructure.

Each attack targets a different layer of the financial services stack. Defending against them **requires a unified view of risk across the full customer lifecycle**, from application through transaction through account maintenance, rather than point solutions at each layer in isolation.


eCommerce, Quick Commerce, and Gig Platforms

- **Real-Time Delivery Creates Fraud Surface:** Ultra-fast fulfillment collapses verification windows, enabling automated address manipulation and fake delivery claims at scale.
- **Organized Networks Drive Return Abuse:** Return fraud has shifted to coordinated rings that exploit policy gaps across merchants, overwhelming manual review teams.
- **Promotional Abuse Inflates Acquisition Costs:** Multi-account farming drains first-user discounts and referral credits, distorting true customer acquisition metrics.
- **Supply-Side Attacks Compromise Gig Platforms:** Fake worker onboarding and payment diversion turn gig accounts into laundering nodes for broader financial crime.
- **eCommerce Checkouts Enable Card Testing:** Attackers exploit checkout flows to validate stolen cards or hijack stored credentials, driving chargeback penalties.

Bureau's fraud intelligence network recorded the share of **high-risk quick commerce sessions nearly doubling over the past year**, with close to 82 million high-risk sessions identified across five quarters. By Q2 2026, more than one in every 23 quick commerce sessions was classified high risk.

Risky events (% of Total Events) by quarter – eCommerce vs. Quick Commerce

Comparative view across quarters (2025 - 2026)



The attacks in these sectors are lower in individual value than financial services fraud, but they operate at a volume that overwhelms manual review. They exploit platform mechanics: return policies, delivery confirmation, promotional structures, and worker onboarding, that were designed for legitimate use with no fraud-resistance built in. And increasingly, these attacks are not individual and opportunistic. They are coordinated, repeatable, and run like operations.

Card Testing and Stored Credential ATO in Checkout Flows



eCommerce checkouts serve a function in the fraud economy that has nothing to do with purchasing goods. They are used to **validate stolen payment card data before it is deployed for higher-value fraud elsewhere**. The technique of running low-value transactions through checkout flows to confirm which card numbers are live enables fraudsters to turn eCommerce platforms into an unwitting testing network for the criminal card market.

The impact is both financial and systemic. Every successful card test generates a chargeback, raises the platform's fraud rate, and can trigger cost penalties or payment processing consequences from card networks at scale. The threat has evolved alongside checkout technology. Stored credential attacks, where an account takeover provides

access to saved payment methods, addresses, and purchase history, represent a higher-value variant of the same attack. The attacker inherits not just a card but trusted account with a clean transaction record, making real-time fraud detection significantly harder.

According to the FTC's April 2026 Data Spotlight⁴⁶, social media has become the primary entry point for eCommerce fraud. Consumers lost \$2.1 billion to these scams in 2025, eight times more than in 2020.

Nearly 30% of victims said the fraud began with contact on a social media platform. Online shopping fraud, driven by fraudulent merchants, fake listings, and non-delivery, has consistently ranked among the most frequently reported fraud categories in FTC consumer data⁴⁷, and social media is now the dominant delivery mechanism for it.

FTC DATA SPOTLIGHT · APRIL 2026

According to the FTC's April 2026 Data Spotlight⁴⁶, social media has become the primary entry point for eCommerce fraud. Consumers lost \$2.1 billion to these scams in 2025, eight times more than in 2020. Nearly 30% of victims said the fraud began with contact on a social media platform. Online shopping fraud, driven by fraudulent merchants, fake listings, and non-delivery, has consistently ranked among the most frequently reported fraud categories in FTC consumer data⁴⁷, and **social media is now the dominant delivery mechanism for it**.

Return and Refund Abuse at Organized, Cross-Merchant Scale

Return fraud is the sector's most underestimated loss category. It does not look like fraud in the conventional sense. This is because the account is legitimate, the original transaction was genuine, and the dispute goes through the platform's own process. That is precisely why it is so difficult to detect and so consistently misclassified as a customer experience problem rather than a fraud loss.

What has changed is who is doing it and how. **Return fraud has moved from individual opportunism to coordinated group activity.** Organized networks systematically identify return policy gaps across multiple merchants, share methods for maximizing claim success, and operate at volumes that no

individual merchant's fraud team is built to catch. The playbooks include returning goods after use, substituting counterfeit items in return packages, claiming non-delivery on orders that were successfully fulfilled, and then filing refunds on orders never legitimately placed.

At the individual transaction level, a fraudulent return is often indistinguishable from a legitimate claim. The diagnostic signal is network-level: the same addresses, the same devices, the same timing patterns appearing across merchants and accounts simultaneously. **That signal only becomes visible with cross-merchant data sharing,** which most platforms do not currently have.



Quick Commerce Fraud: Address Manipulation, Fake Delivery Claims, and Promo Abuse

Quick commerce has grown rapidly across India, the UK, Europe, and Southeast Asia. With deliveries promised in as little as 10 to 30 minutes, it has created a new fraud attack surface that did not exist five years ago. The vulnerabilities follow directly from how the model works: rapid fulfillment, no time for pre-delivery verification, heavy promotional spend to drive adoption, and a delivery confirmation process that relies largely on self-reporting. Here's how fraud plays out:

ADDRESS MANIPULATION



Exploits the gap between order speed and fulfillment accuracy. Orders placed with modified or boundary-straddling addresses route deliveries to unintended locations while leaving the platform unable to confirm non-delivery. There is no time to verify before dispatch.

FAKE DELIVERY CLAIMS



Exploit how platforms resolve disputes. In high-density urban markets, where quick commerce operates, contested delivery claims are common and expensive to investigate. Platforms typically resolve them in the customer's favor to protect satisfaction scores. Organized groups exploit that resolution pattern systematically.

PROMOTIONAL ABUSE




Is the highest-volume attack. Quick commerce platforms spend heavily on first-order discounts, referral credits, and free delivery thresholds to acquire users. Account farming, which entails creating multiple accounts to claim first-user promotions repeatedly, uses the same process as synthetic identity construction. At scale, it can materially distort a platform's cost-per-acquisition without appearing in any fraud loss figure.

Bureau's fraud intelligence network shows **suspicious multi-device behavior tied to promotional abuse rose sixfold in quick commerce**, while eCommerce held broadly flat over the same period. The activity is also geographically concentrated rather than evenly distributed, consistent with coordinated multi-account operations rather than isolated coupon misuse.

Promo Abuse: Users with more than 4 Devices

Suspicious user rate (%) trend across industries & quarters.



KEY TAKEAWAYS:

2025

Promo abuse risk increased across the year in Quick Commerce, rising from 1.49% in Q2 to 5% in Q4. eCommerce also saw an increase in Q3 and peaked at 15% in Q4

2026

Quick Commerce saw a sharp drop from 60% in Q1 to 10% in Q2 2026. E-commerce declined significantly from 9% in Q1 to 0.5% in Q2 2026

Gig Platform Identity Fraud: Fake Worker Onboarding, Payment Diversion, and Rating Manipulation

Gig platforms face a different fraud profile from eCommerce: the attack is on the supply side, not the demand side. The accounts being created or compromised are worker accounts with different verification requirements, payment flows, and consequences when they go wrong.

Uses synthetic and fabricated identities to create accounts where verification standards are often less rigorous than financial services KYC. Once active, these accounts can divert payments to fraud-controlled bank accounts, inflate earnings through coordinated self-ordering schemes, or operate as mule accounts, linking gig platforms to the broader financial crime ecosystem.	Typically follows social engineering of a legitimate worker account — an impersonation of platform support that redirects payment details to a fraudster-controlled account. The FBI's November 2025 PSA on account takeover fraud ⁴⁹ documented the same playbook operating across financial institutions, employment platforms, and payment services alike.	Coordinated fake reviews, competitor downgrading, and inflated completion records degrade the trust infrastructure that marketplace platforms rely on to allocate work and protect consumers. The direct financial loss is lower; the damage to platform integrity is not.
---	--	--

Marketplace Seller Fraud: Fake Merchants, Counterfeit Goods

Third-party marketplace models, where independent sellers transact through a platform that provides payment processing, trust infrastructure, and consumer recourse, create a fraud attack surface that grows with the platform itself.

The pattern is consistent across markets. A seller account passes initial verification, builds a positive rating through legitimate low-value transactions, then scales to high-value fraud: shipping counterfeit goods, collecting payment for items never dispatched, or using the merchant account as a layer in a payment laundering chain. **The account is flagged and closed, but the process has already paid out.** Europol's IOCTA 2026⁵⁰ identifies marketplace fraud as part of a broader shift to increasingly fragmented online fraud channels that resist platform-level enforcement alone. Escrow abuse exploits the consumer protection mechanisms platforms introduce to build trust. Fraudulent sellers manipulate the timing of delivery confirmation, dispute resolution, and escrow release to extract payment before a dispute can be raised. This is particularly true for high-value categories where buyers are motivated to complete quickly and verification is slow.

BUREAU CLOSE

Every fraud vector, explained above, exploits a feature designed for legitimate use. Return policies, delivery confirmation, promotions, and worker onboarding all rely on trust. **Without a fraud intelligence layer, they become attack surfaces.** The platforms that manage fraud most effectively build detection into the product itself, rather than relying on reviews after the transaction is complete.

SECTION 02

The Economic and Regulatory Impact

CH 07

The True Cost of Fraud:
Beyond the Loss Line

CH 08

The Global Compliance
Landscape

CH 09

The Cost of
Non-Compliance

The True Cost of Fraud: Beyond the Loss Line

- **Direct Losses Mask the 6-Layer Iceberg:** Write-offs capture only Layer 1, while prevention costs, friction, brand erosion, regulatory liabilities, and opportunity costs remain hidden.
- **Prevention Spend Outstrips Direct Losses:** Investigations, case reviews, and operational overhead can cost more than the fraud itself.
- **Customer Friction Harms Revenue:** False declines and drop-offs turn fraud controls into hidden revenue loss.
- **Regulatory Penalties Become Fixed Costs:** APP reimbursements and regulatory fines turn compliance into an ongoing operating expense.
- **Reactive Defense Taxes Strategic Growth:** Fraud mitigation diverts capital and leadership attention from innovation and expansion.

The Fraud Iceberg

LAYER	WHAT IT COVERS	VISIBLE TO THE BOARD?
Direct Fraud Losses	Confirmed losses net of recoveries	
Prevention and Operational Response	Detection infrastructure, investigation, claims handling	RARELY
Friction, False Declines and Abandonment	Revenue lost to over-detection	NO
Brand Depreciation and Trust Erosion	LTV reduction, acquisition drag post-incident	NO
Regulatory Penalties and Remediation	Fines, mandatory reimbursement, compliance spend	WHEN ENFORCED
Strategic Opportunity Cost	Growth capital diverted to reactive response	NEVER

Direct Fraud Losses: The Floor, Not the Ceiling

US consumers reported \$15.9 billion in fraud losses to the FTC in 2025⁹, up 27% from 2024, and 430% above 2020 levels. The FTC's own estimate, accounting for underreporting, puts actual US consumer fraud losses for 2024 as high as \$195.9 billion. That is the distance between what reporting systems capture and what is actually happening.

The pattern holds globally. In the UK, payment fraud losses reached £1.28 billion in 2025¹⁰, the fourth consecutive year above £1 billion. In India, reported bank fraud reached ₹48,021 crore in FY2025-26¹¹.

These figures are the starting point for understanding what fraud costs. However, they are not the answer.



LAYER 02

Prevention and Operational Response: The Cost That Exceeds the Loss

In mature markets, **what institutions spend preventing fraud now routinely exceeds what they lose to it**. This is not a sign of success. It is a sign of how resource-intensive the containment problem has become; and how little of that cost appears in fraud reporting.

The scale becomes clear at the operational level. Bureau's fraud intelligence network processed more than 3.2 billion eCommerce events in Q2 2026 alone. Just 0.76% were classified as high risk, yet that still amounted to 24.6 million events requiring review, more than double the 11.2 million flagged in the prior quarter. Even a fraud rate below 1% can generate millions of cases, overwhelming fraud teams with investigations, reviews, and false positive resolution quarter over quarter.

Prevention has a measurable impact, but only on fraud that systems can see. UK financial institutions blocked £1.68 billion in unauthorized fraud attempts in 2025, stopping 70 pence of every £1 flagged as suspicious. Yet £703.4 million still resulted in financial losses¹². The figures highlight an important limitation: prevention rates measure what institutions detect and act on, not the fraud that remains invisible to their systems. For APP fraud, institutions processed 306,000 mandatory reimbursement claims worth £215 million, against £576.4 million in reported APP fraud losses¹³.

Singapore's Anti-Scam Command recovered S\$140.5 million in scam proceeds in 2025 and averted a further S\$348 million in potential losses, an operational infrastructure that exists entirely outside the S\$913.1 million loss figure it contains¹⁴.

In India, the RBI deployed MuleHunter.AI across 23 banks, detecting approximately 20,000 mule accounts per month¹⁵. The incorporation of the Indian Digital Payment Intelligence Corporation (IDPIC) in October 2025 further signals that fraud prevention is moving toward shared,

real-time intelligence across the digital payments ecosystem, not isolated defenses at individual institutions. Prevention spend is the most visible of the hidden layers. Yet it almost never appears on the same report as the losses it contains..

UK FINANCE 2026

Friction, False Declines and Abandonment: The Revenue Problem Inside the Fraud Problem

Every fraud control that stops a fraudster also stops some legitimate customers. That revenue cost is real. It lives in product analytics, not fraud reports, which is precisely why it disappears from the investment conversation.

FTC estimates that only 4.8% of mass-market fraud victims report their losses to a government entity¹⁶. The internal complaint, dispute, and review volumes that precede and surround each reported case, and the friction experienced by the legitimate customers caught in the same controls, **represent a customer cost that finds no place in the loss data.**

False declines and abandonment are a revenue problem dressed as an operational one. They are invisible to the teams making fraud investment decisions because they are reported, if and when they get reported at all, to the wrong function.



FTC · REPORTING RATE

4.8%

of mass-market fraud victims report their losses to a government entity

LAYER 04

Brand Depreciation and Trust Erosion: The Long Tail of a Fraud Event

Fraud events redistribute trust. This redistribution lasts longer than the incident. Institutions bear reputational costs not based on their technical liability for the fraud, but on how customers experience the response.

Of the UK APP fraud victims who went through the mandatory reimbursement process, 42% said their

trust in their bank increased¹⁷. But around one in three of those who were not reimbursed, also reported the same. The data makes the implication structural: how an institution handles fraud, not just whether it prevents it, determines the trust outcome. And **trust outcomes determine customer lifetime value** in ways that no direct loss figure captures.

Regulatory Penalties and Remediation: A Scheduled Cost, Not a Tail Risk

Across the UK, EU, and Australia, regulatory cost has moved from contingent liability to fixed operating expense. It scales with fraud volume, grows as enforcement matures; and unlike direct losses, carries personal accountability for leadership.

European data protection authorities issued approximately €1.15 billion in GDPR fines in 2025¹⁸, bringing the cumulative total since 2018 to €7.1 billion. The EU AI Act adds a second layer: fraud detection, credit scoring, and identity verification systems are classified as high-risk AI under Annex III, with penalties for non-compliance of up to €15 million or 3% of global annual turnover.

In the UK, PSPs reimbursed £215 million to APP fraud victims in 2025. In the first 15 months of the mandatory

reimbursement regime, 89% of money lost to APP scams, £243 million, was returned to victims¹⁹. This is a regulatory liability that did not exist before October 2024, and that scales directly with fraud volume.

In Australia, the Scams Prevention Framework Act 2025²⁰ imposes civil penalties of up to AUD \$50 million per contravention on banks, telecoms, and digital platforms, with full obligations taking effect March 31, 2027.

In India, an RBI consultation on compensation for unauthorized electronic banking transactions²¹ signals mandatory victim liability approaching.

The direction is consistent across every major operating jurisdiction. Regulatory cost is becoming mandatory, measurable, and enforceable.

LAYER 06

Strategic Opportunity Cost: The Growth Tax

The least visible layer carries the most compounding effect. Capital and leadership attention directed at reactive fraud response are not available for product development, market expansion, or customer experience investment. This cost never appears on any ledger; it appears as decisions not made, markets entered more slowly, and roadmaps that run perpetually behind.

According to FTC's March 2026 data²², US consumer fraud losses rose 430% between 2020 and 2025. Every institution operating across that trajectory has had to continuously scale reactive fraud response, operations, technology, compliance, legal, at multiples of underlying business growth. The capital absorbed by that scaling is a tax on the growth agenda, charged annually and compounding silently.

Building the Board Case

Instead of understanding that the six layers exist, the real challenge for executive teams is making them visible enough to change investment decisions.

Three steps that can help close that gap are:

Quantify what is already tracked.

Direct losses plus prevention spend. In most institutions, this alone reveals that fraud's institutional cost is substantially larger than the reported loss figure: before friction, brand, or regulatory exposure are counted.

STEP 02

Estimate what can be measured.

False decline rates and post-decline churn sit in product analytics. One quarter of that data, mapped against the fraud loss figure, typically reveals a revenue cost of comparable magnitude to direct losses.

Build the regulatory liability line.

GDPR penalties, APP reimbursement exposure, SPF obligations, and emerging compensation mandates are quantifiable in advance of enforcement. Model them as scheduled costs, not contingent risks.

The board case is not: fraud costs \$X. It is: tracked \$X of a cost that is structurally larger; and the gap between what tracked and what was carried is where investment decisions are being made blind.

BUREAU CLOSE

Fraud economics have changed. Institutions that build their fraud investment case around direct losses alone are solving for the visible fraction of a larger cost, and leaving the rest to accumulate invisibly across prevention spend, customer friction, regulatory liability, and strategic diversion. Closing that gap starts with measuring the right number.

The Global Compliance Landscape

- **Regulatory Mandates Shift Loss Liability to Institutions:** Frameworks like the UK's PSR, Australia's SPF, and EU's PSD3/PSR turn fraud into a direct institutional financial liability.
- **Global Convergence Closes Regulatory Arbitrage Gaps:** Synchronized AML, KYC, and conduct standards across regions reduce the regulatory blind spots exploited by transnational criminal networks.
- **Governed AI Replaces Black-Box Models:** Compliance regimes like the EU AI Act require transparent, auditable, and explainable AI decisioning alongside detection performance.
- **Cross-Border Intelligence Combats Networked Fraud:** Initiatives like INTERPOL's I-GRIP, Singapore's Anti-Scam Command, and India's IDPIC address the critical need for shared fraud data.
- **Privacy Laws Constrain Global Data Architectures:** Tensions between cross-border fraud intelligence and laws like GDPR or DPDP drive adoption of Privacy-Enhancing Technologies (PETs).

For decades, financial crime regulation evolved in silos, with inconsistent enforcement and significant divergence in compliance standards. Criminal networks exploited those gaps, routing operations through weaker jurisdictions, exploiting data-sharing blind spots, and turning regulatory fragmentation into an advantage. That is changing. Across the US, UK, EU, Australia, India, Singapore, and the GCC, regulatory priorities are converging around AI accountability, greater fraud liability, cross-border collaboration, and stricter conduct standards. Differences remain, but regulatory arbitrage is becoming harder, and the compliance bar is rising globally.

The Synchronization That Matters

The significance of this regulatory shift is not any single framework, but the convergence between them. Institutions can no longer manage fraud compliance as a series of separate, jurisdiction-specific obligations. A more consistent global approach is emerging.

That approach rests on four pillars.

Govern: AI system accountability

As AI becomes central to fraud detection and decisioning, regulators are scrutinizing how these systems make and govern decisions, not just their outcomes. The EU AI Act is setting an early benchmark, with requirements around risk management, transparency, human oversight, and accountability for high-risk AI systems. For fraud teams, this means AI can no longer operate as a black box: institutions need governed, explainable, and auditable decisioning alongside detection performance.

Prevent: Mandatory liability

The shift from voluntary best practice to enforceable obligation is the defining characteristic of the current regulatory generation. As fraud becomes more sophisticated, cross-platform, and difficult to prevent, regulators are placing greater responsibility on institutions that have the infrastructure, data, and controls to detect and intervene. The UK's PSR APP reimbursement mandate, Australia's Scams Prevention Framework, India's unauthorized transaction compensation framework, and the EU's PSD3/PSR liability provisions all reflect the same principle: institutions that process payments increasingly bear a defined share of the fraud losses occurring within their infrastructure. The economic consequence is material and growing, for example, UK PSPs reimbursed £243 million to APP fraud victims in the first 15 months of the mandatory regime.

Detect: Intelligence Must Cross Boundaries

Regulators are recognizing that no institution can fight increasingly coordinated fraud alone. They are promoting cross-institutional intelligence sharing and centralized fraud data through initiatives such as Singapore's Anti-Scam Command, India's Indian Digital Payment Intelligence Corporation, and the UK's National Fraud Database, while INTERPOL's I-GRIP extends collaboration across borders.

Together, these efforts support the data sharing required to detect fraud across networks rather than within individual organizations. The opportunity is visible at the institutional level in fast-growing markets, too. Bureau's research among senior fraud and risk leaders across the Middle East found **there's a meaningful room to grow in cross-institutional collaboration**, with most describing their information-sharing efforts as still limited or just beginning.

The Data Sovereignty Tension


Industrialized fraud does not respect jurisdictional boundaries, but intelligence must. Attacks can originate in one country, target victims in another, and reuse identities or devices across several markets. Real-time data sharing between countries is, therefore, critical to connecting signals that individual institutions cannot see alone. For example, linking a device flagged in Singapore to a fraud cluster in Indonesia, or an identity presented for KYC in the UAE to a prior rejection by a UK institution. The same fragmentation exists within countries, where state, provincial, and institutional boundaries can make even domestic networks difficult to investigate.

The regulatory environment creates a direct tension with this requirement. GDPR's extraterritoriality provisions, India's DPDP Act data localization requirements, and China's data export controls constrain the data architectures that cross-border fraud intelligence needs. The challenge is therefore not whether institutions need to share intelligence, but how they can do so within privacy and data-sovereignty boundaries. Building compliant mechanisms for cross-border fraud intelligence is becoming essential to keeping pace with fraud networks that already operate without borders.

Three architectural approaches are emerging to enable secure fraud intelligence sharing.

Privacy-enhancing technologies

such as federated learning, differential privacy, and secure multi-party computation, **are moving from concept to practical use**. The UK is developing PET-based financial-crime use cases through public-private initiatives, Australia's AUSTRAC is building a PET-enabled platform to identify suspicious links across financial institutions, and Singapore's IMDA has established a PET Sandbox, including a proof of concept for sharing financial-crime intelligence across Singapore, the US, India, and the UK.



Tokenization and pseudonymization

protect individual identities while preserving the value of fraud signals for analysis.

Industry consortium models

bring institutions together under shared legal and governance frameworks, enabling cross-institutional data sharing that meets privacy requirements and provides visibility no single organization can achieve alone.

None of these is a complete solution. But **each addresses a real constraint**, and the institutions investing in them now are building the data architecture that the next generation of cross-border fraud regulation will reward.

PILLAR 04

Respond: Liability is Becoming Predictable.

Regulation is defining not only what institutions must prevent, but what happens when prevention fails. Mandatory reimbursement, reporting obligations, investigation requirements, and regulatory penalties are turning fraud response into a defined institutional responsibility. The economic consequence is already visible. As liability becomes defined by regulation rather than negotiated case by case, institutions face a more predictable cost when fraud controls fail.

The Jurisdiction Map

The table below reflects the regulatory development and enforcement intensity across countries:


JURISDICTION	KEY FRAMEWORKS ACTIVE OR APPROACHING	ENFORCEMENT INTENSITY
United States	FedNow expansion; P2P liability framework evolving at state level; federal AI guidance forthcoming	— federal enforcement in transition; state-level escalating
United Kingdom	PSR mandatory APP reimbursement (October 2024); FCA AI expectations developing; parallel operational resilience regime under PS24/16 (critical third-party oversight, since January 2025)	— mandatory reimbursement creates direct P and L liability
European Union	PSD3/PSR (entry into force 2027); EU AI Act (high-risk obligations August 2026 / December 2027); DORA (January 2025); MiCA (December 2024)	— most comprehensive and enforceable stack globally
Singapore	MAS Shared Responsibility Framework; Facility Restriction Framework (October 2025); Digital Token guidelines	— MAS among the most proactive regulators globally
India	DPDP Act; RBI digital lending guidelines; unauthorized transaction compensation consultation (March 2026)	— framework maturing; implementation pace accelerating
Australia	Scams Prevention Framework Act 2025 (full obligations March 2027); AFCA complaint resolution mandates	— tri-regulator enforcement model operational
UAE / Saudi Arabia	CBUAE OTP elimination (March 2026); SAMA Counter-Fraud Fundamental Requirements (April 2026)	— rapidly maturing; GCC enforcement sophistication increasing

The End of Regulatory Arbitrage

The practical consequence of regulatory synchronization is straightforward: regulatory arbitrage is becoming less viable. As jurisdictions strengthen AML, KYC, and financial crime controls, the variation in regulatory standards that once enabled networks to exploit weaker environments is narrowing.

The UAE illustrates this shift. After sustained investment in financial crime controls, it was removed from the FATF grey list in February 2024 and the EU high-risk list in July 2025, closing a route once available to illicit flows. Similar tightening is underway across the GCC, Southeast Asia, and digital assets under MiCA.

For fraud networks, the implications are clear: **higher operating costs, greater exposure, and fewer regulatory differences to exploit.**



BUREAU CLOSE

The global regulatory environment is converging on **a consistent set of expectations: mandatory liability, explainable AI, cross-border data collaboration, and enforceable conduct standards.** FATF ministers signaled the same direction in their April 2026 Ministerial Declaration, committing jointly to combat fraud “in all its forms,” including organized scam centers and fraud enabled by emerging technologies such as AI¹⁰⁵.

For institutions, a jurisdiction-by-jurisdiction approach will struggle to keep pace with a regulatory landscape that is becoming more coordinated and demanding. The institutions getting ahead are **treating this direction as a strategic signal; building for where regulation is headed,** rather than waiting for each requirement to become mandatory.

The Cost of Non-Compliance

- **Fraud Non-Compliance Is a Forecastable Liability:** Mandates like UK APP reimbursement turn fraud exposure into a calculable balance-sheet cost rather than a contingent risk.
- **Severe Penalties Include Operational Limits:** Regulators enforce strict consequences, ranging from massive EU AI Act fines to restrictions on onboarding new customers.
- **Proactive Defense Delivers 21:1 ROI:** Upstream fraud prevention yields an estimated 21:1 return on investment, far outperforming reactive remediation.
- **Ad-Hoc Remediation Inflates Overall Costs:** Emergency fixes under regulatory scrutiny drive up implementation fees, limit architectural flexibility, and damage reputation.
- **Board Governance Must Quantify Financial Risk:** Boards must track expected reimbursement liability, AI compliance gaps, and potential enforcement penalties as core enterprise risks.

Regulatory risk used to be a tail risk; something that affected other institutions after a visible failure. Today, it is a direct business exposure. Weak fraud controls can create six interconnected costs: regulatory penalties, fraud losses, reimbursement liability, operational costs, customer impact, and reputational damage.

Across the UK, EU, Australia, and India, fraud-related regulatory exposure is increasingly calculable in advance. Liability can scale with payment volumes, fraud incidents, and AI systems deployed today. **Enforcement is also moving beyond fines:** there are real-world instances where regulators have restricted institutions from onboarding new customers, required them to suspend specific products or activities, and imposed limits until control deficiencies were addressed.

A control failure can therefore trigger penalties, reimbursement, investigation and remediation costs, customer friction, and reputational damage; far beyond the initial fraud loss. For institutions operating under mandatory reimbursement, AI governance, or cross-sector scam prevention requirements, regulatory liability is already accumulating, whether it has been modeled or not.

Regulatory risk is now a measurable component of fraud risk. **Stronger fraud infrastructure reduces exposure across all six cost categories**, and helps protect not just the balance sheet, but the institution's ability to operate.

From Reputational Risk to Scheduled Liability

Fraud is moving from a risk institutions absorb to a liability they are expected to manage. Regulators are making the consequences of weak fraud controls increasingly predictable: obligations have defined effective dates, liability is assigned by rule, and penalties or reimbursement costs can be quantified in advance. That **predictability is changing how financial institutions approach fraud prevention**: from containing losses after they occur to investing upfront in controls that reduce their regulatory exposure.

The UK's mandatory APP reimbursement regime¹⁰⁰ illustrates the shift. Since October 2024, reimbursement has stopped being a matter of institutional discretion. £243 million moved from PSPs to victims in the regime's first 15 months alone, not because institutions chose to pay it, but because a rule with a defined start date and a defined cost structure required them to.

Australia's Scams Prevention Framework Act 2025¹⁰¹ extends the principle across banks, telecoms, and digital platforms, with civil penalties of up to AUD \$50 million per contravention for failures to meet anti-scam obligations.

The EU AI Act¹⁰² adds another dimension: accountability now extends to the decision-making system itself, not just its outcome, with exposure reaching €15 million or 3% of global turnover.

The paradigm shift is clear: fraud prevention is no longer simply a question of balancing fraud losses against operational cost. The cost of inadequate prevention is becoming a defined and recurring regulatory exposure.

For financial institutions, investing in stronger detection, prevention, and governance is a way to manage a predictable liability, not simply reduce an unpredictable loss.

The Cost Structure of Non-Compliance

The regulatory cost of inadequate fraud infrastructure now operates through three increasingly predictable mechanisms: reimbursement liability, conduct penalties, and AI governance obligations. But the bigger shift is in how institutions approach fraud itself. Reactive compliance: responding after losses, regulatory findings, or enforcement, is becoming expensive. The frontrunners are moving upstream, using intelligence, technology, and continuous controls to identify and address risk before it becomes a loss or a regulatory breach.

Yet preparedness remains uneven. Many institutions are underprepared for emerging fraud threats, particularly AI-enabled attacks, while leading organizations are moving toward more proactive, continuous approaches to fraud prevention. The result is a widening gap between institutions that respond to fraud as an event and those that manage it as an ongoing risk.

Scales with fraud volume and payment rail activity. For institutions operating on UK Faster Payments, EU instant credit transfer, or Indian UPI rails, annual exposure can increasingly be modeled from fraud rates, transaction volumes, and statutory reimbursement requirements. What was once a contingent risk is becoming a forecastable financial liability.

Arise when institutions fail to meet defined obligations, such as reporting confirmed fraud actors or implementing required anti-scam controls. The FCA's December 2025 review found that 63% of confirmed money mules were not reported to the National Fraud Database, with the FCA disagreeing with firms' decisions in nearly one-third of cases.

adds another layer of accountability. EU AI Act requirements covering documentation, risk management, bias testing, explainability, and human oversight make AI governance an ongoing obligation, not a one-time implementation exercise. GDPR enforcement¹⁰³ provides a precedent: EU data protection fines reached €7.1 billion between 2018 and 2025, with enforcement expanding beyond technology companies into regulated industries.

The dividing line is clear: reactive institutions absorb fraud and regulatory costs as they arise; proactive institutions invest ahead of them. By moving upstream, frontrunners are reducing losses, improving operational efficiency, strengthening compliance readiness, and building more resilient fraud defenses.

The economics favor prevention over response. An analysis by the UK Public Sector Fraud Authority found approximately 21:1 ROI for proactive counter-fraud activity, compared with around 5:1 for reactive measures¹⁰⁶.

The principle is clear for financial institutions: the earlier fraud risk is addressed, the more costs can be avoided across investigation, recovery, remediation, and regulatory response. Prevention is not simply a cost of doing business; **it is an investment whose value compounds when risk is addressed before losses occur.**

Remediation happens under time pressure, which drives up vendor costs, implementation risk, and the likelihood of building to the minimum required rather than to a defensible standard. It happens under regulatory scrutiny, which constrains design choices and requires continuous evidence production alongside the remediation work.

And it happens after the reputational event: enforcement action, publicized fraud failure, or the regulatory letter, that changes how customers, counterparties, and regulators view the institution.

PROACTIVE INVESTMENT

Proactive investment happens on the institution's timeline, to the institution's specification, with the full range of architectural choices available. The infrastructure built proactively also tends to be better, designed for the institution's specific risk profile rather than constructed to satisfy an immediate regulatory requirement.

Proactive investment in fraud and compliance infrastructure creates a competitive advantage. Institutions with strong data-sharing, AI governance, and fraud intelligence foundations face lower compliance costs and can enter new markets faster.

THE BOARD CONVERSATION

The regulatory liability landscape requires a different board conversation. Most institutions still frame fraud compliance around costs: annual spend, direct fraud losses, and next year's budget. This treats compliance as an expense rather than a measurable liability.

The better approach is to quantify exposure: expected reimbursement liability, AI governance gaps, potential conduct findings, and the financial impact of regulatory risk. These are measurable risks that belong on the enterprise risk register, not only in compliance reporting.

BUREAU CLOSE

The **shift from voluntary best practice to mandatory, penalized, and automated regulatory enforcement has happened.** Institutions still treating fraud compliance as a reputational and legal risk rather than a financial one with a calculable cost structure, are not just behind on compliance. They are behind on risk management.



SECTION 03

Regional Risk Intelligence



CH 10	CH 11	CH 12	CH 13	CH 14
The United States	Europe and United Kingdom	Southeast Asia	Asia-Pacific	Middle East and North Africa

The United States

- **Record Scale Outpaces US Detection Infrastructure:** Daily complaint volumes push total internet crime losses to \$20.877 billion, the highest figure in the FBI IC3's 25-year history⁵¹.
- **FedNow Limit Hikes Increase Exposure:** Raising FedNow limits to \$10 million on instant, irrevocable rails creates pre-authorization risk.
- **State Enforcement Governs P2P Regulation:** Regulatory accountability for authorized push payment fraud is migrating from federal oversight to state attorneys general.
- **BEC and AI Fraud Drive Losses:** Business Email Compromise leads corporate financial losses, while tracked AI-facilitated fraud reached \$893 million.
- **Community Banks and BNPL Carry Highest Vulnerability:** Regional banks using instant rails and unregulated BNPL providers face synthetic identity and pre-authorization exposure.



Three characteristics define the US fraud environment in 2026.

Scale outpaces detection.	At nearly 3,000 fraud complaints per day to the FBI IC3 alone, the volume of reported fraud has crossed a threshold where manual review and case-by-case investigation are structurally insufficient. The average reported loss per complaint reached \$20,699 in 2025, the highest on record ⁵² . The pattern extends beyond US-reported crime data. Bureau's fraud intelligence network detected approximately 11.6 million suspected phishing attempts across 7.3 billion digital sessions globally, a rate of just 0.16%, yet still translating into millions of high-risk interactions capable of fueling downstream account takeover and payment fraud.
Real-time payment infrastructure is expanding faster than fraud controls	FedNow has grown to over 1,500 participating financial institutions across all 50 states. This is a 44% increase in 2025. In November 2025, the Federal Reserve raised FedNow's per-transaction limit from \$1 million to \$10 million, opening the network to high-value commercial use cases ⁵³ . Higher transaction limits on an instant, irrevocable rail without a commensurate uplift in pre-authorization fraud intelligence is a risk multiplier, not just a product upgrade.
The P2P liability framework is evolving	A landmark federal enforcement action in December 2024 cited \$870 million in consumer losses on peer-to-peer payment networks since 2017, establishing the scale of the problem on the public record. A subsequent state-level action in August 2025 cited over \$1 billion in losses to residents of a single state over the same period. Consumer protection enforcement in the P2P space is now being actively pursued at the state level, with several states advancing their own digital payment legislation. The result is a national framework in active development : one that institutions operating across state lines need to monitor closely.

Three characteristics define the US fraud environment in 2026.

Business Email Compromise remains the highest dollar-loss fraud category targeting US businesses. In 2025, the FBI IC3 recorded 24,768 BEC complaints with losses of ~\$3.1 billion, an increase from \$2.77 billion in 2024, and the second-highest loss category behind investment fraud⁵⁴. BEC requires no technical compromise. It **exploits payment authorization processes** through impersonation, targeting the human decision-maker, not the system.

~\$3.1 Billion

Losses in 2025

~25,000

Complaints

Synthetic identity fraud in consumer credit and BNPL is the silent loss driver. The BNPL market reached an estimated \$70 billion in transaction value in 2025⁵⁵, with 61% of BNPL users carrying subprime or deep subprime credit scores, the same thin-file population that synthetic identity construction specifically targets. The regulatory transition in BNPL oversight means **fraud exposure in this segment is growing** at a moment when the compliance framework is still being defined.

\$70 billion

Transaction value in 2025

61%

BNPL users carrying subprime or deep subprime credit scores

P2P payment fraud remains structurally unresolved. The debate has moved beyond whether fraud is happening to who pays for it. Instant P2P payments create a unique challenge. Because **victims authorize the payment through social engineering** rather than technical compromise, the losses fall into the regulatory gray area between authorized and unauthorized transactions. That distinction is now **the subject of active state-level litigation and legislative attention**. Its resolution will determine the liability architecture for the next decade of consumer payments

AI-facilitated fraud was formally tracked as a standalone category in the US for the first time in 2025. The FBI IC3 recorded 22,364 AI-related complaints with \$893 million in attributed losses⁵⁶. It noted that both figures are undercounts because most victims do not recognize AI involvement in the attacks targeting them.

\$893 Million

Attributed loss

22K+

AI-related complaints

Regulatory Pulse

The US regulatory landscape is one of active transition, with responsibility for consumer fraud protection shifting meaningfully from federal to state level.

State attorneys general are taking a more active role in regulating P2P payment fraud. Several states are advancing fraud liability laws and enforcement actions that will influence how institutions manage authorized push payment (APP) risk. As federal oversight shifts toward greater industry self-regulation, state-level action is shaping the accountability framework for both consumers and financial institutions.

On AI-specific fraud guidance, the FATF's February 2026 Cyber-Enabled Fraud paper⁵⁷, applicable to US institutions under existing AML and consumer protection frameworks, **provides the clearest current framework for treating AI-generated fraud** as an institutional risk requiring a documented response. Dedicated US federal AI fraud guidance is expected to follow as the threat profile becomes more visible in reporting data.

Industry Hotspots

Community and regional banks


Community and regional banks carry disproportionate exposure on the FedNow rail. Larger banks built fraud detection infrastructure ahead of real-time payment adoption. Community institutions, which account for a significant portion of FedNow's 1,500+ **participating institutions, are implementing instant payment capability without equivalent fraud prevention investment.**

BNPL and embedded finance

BNPL and embedded finance providers operate at the junction of rapid credit extension and an evolving regulatory framework. With BNPL oversight still being defined at the federal level, **fraud programs in this segment need to be built ahead of regulation** rather than in response to it.

Corporate treasury teams

Corporate treasury teams face growing BEC risk as FedNow's higher transaction limits enable instant, irrevocable transfers of larger payment amounts. **Protecting instant transfers** of up to \$10 million **requires a different level of pre-authorization risk assessment** than was needed when transaction limits were \$1 million.



BUREAU CLOSE

The **US market rewards speed and punishes friction**. This means fraud programs face a harder trade-off here than in markets with stronger mandatory liability frameworks. In the absence of a federal reimbursement mandate equivalent to those now operating in the UK and Australia, the incentive structure for investment in pre-authorization fraud intelligence is weaker than the threat level requires.

The institutions getting ahead of this will be the ones building pre-authorization risk intelligence into the payment flow: treating every high-value instant transfer as a decision point that requires real-time behavioral and network signal assessment before funds leave the institution, not a review that happens after they have cleared.

Europe and United Kingdom

- **UK Reimbursements Shift Liability:** Mandatory APP rules forced payment providers to return £243 million to victims, turning fraud into a direct operating expense.
- **APP Attacks Extract Higher Value:** UK APP fraud losses jumped 19% despite lower case volumes, driven by AI voice cloning and targeted social engineering.
- **EU Losses Reach Systemic Scale:** Member states suffered \$64.1 billion in fraud losses, led by cross-border crypto scams, online schemes, and BEC.
- **Post-Brexit Rules Create Regulatory Seams:** Divergent UK and EU data-sharing policies leave border gaps that criminal networks exploit.
- **EU Frameworks Impose Strict Compliance:** Approaching PSD3/PSR rules, DORA resilience standards, and EU AI Act penalties mandate heavy operational compliance.

Regional Threat Signature

£576.4M

LOSSES • UP 19% ON THE PRIOR YEAR

Authorized push payment fraud, where victims are deceived into authorizing transfers to fraudster-controlled accounts, reached £576.4 million in 2025, up 19% on the prior year . This was when UK’s financial institutions simultaneously prevented £1.68 billion in unauthorized fraud and reimbursed £215 million to APP victims under the mandatory reimbursement regime⁵⁹. **The architecture of APP fraud: instant, irrevocable, socially engineered**, is specifically designed to complete before any intervention is possible. Higher case volumes would be cause for concern; the fact that APP losses rose 19% while case volumes fell 8% means each individual attack is extracting more⁶⁰.

£1.68B

UNAUTHORIZED FRAUD PREVENTED

£215M

REIMBURSED TO APP VICTIMS

THE EU FACES FRAUD AT SYSTEMIC SCALE

\$64.1B

FRAUD LOSSES ACROSS MEMBER STATES, 2025

Europol's IOCTA 2026⁶¹ recorded \$64.1 billion in fraud losses across EU member states in 2025, naming fraud the fastest-growing area of organized crime in the bloc. Online fraud schemes, investment fraud, and BEC collectively account for the dominant share of reported losses, with AI-assisted social engineering and cross-border crypto laundering as the primary operational accelerants.

THE POST-BREXIT LANDSCAPE HAS CREATED NEW OPPORTUNITIES FOR ORGANIZED FRAUD

UK and EU authorities now operate under separate data-sharing protocols, different investigative frameworks, and diverging financial services regulations. **Cross-border fraud operations route across that seam by design**, placing infrastructure on whichever side of the boundary has the weaker controls at each step in the fraud chain.

Top Attack Vectors

APP fraud and social engineering

APP fraud and social engineering dominate the UK threat landscape across all channels: phone, SMS, email, and social media. These attacks have evolved significantly. Fraudsters now impersonate bank security teams, HMRC, and law enforcement using AI-powered voice cloning and real-time scripting. As a result, many of the cues that once helped victims recognize fraudulent calls and messages have been eliminated. Europol's EU-SOCTA 2025⁶² documented voice cloning and live video deepfakes as operational tools deployed at scale in CEO fraud and romance scams across both UK and EU markets.

Channel-level data highlights where social engineering is most effective. Bureau's fraud intelligence network found that in-app browser sessions consistently carried higher fraud rates than standard browser sessions, which averaged just 0.13% across nearly 250 million sessions. This points to authentication journeys within social and messaging platforms as a higher-risk channel.

Investment fraud

Investment fraud is the highest dollar-loss category across the EU, with cross-border crypto investment scams targeting retail investors across all member states through AI-personalized social media advertising. The combination of seamless digital distribution and irreversible crypto settlement makes recovery functionally impossible in most cases.

Identity fraud and account takeover

Identity fraud and account takeover continue to grow in the UK, with Cifas Fraudscape 2026⁶³ recording identity fraud as the dominant fraud-risk category. It accounted for the majority of all cases filed to the National Fraud Database in 2025, driven increasingly by AI-generated documentation and deepfake-assisted KYC bypass.

Money mule networks

Money mule networks operating across the UK-EU boundary represent a specific cross-border challenge. Mule accounts recruited in one jurisdiction route funds through another before extraction, exploiting the gap in real-time cross-border data sharing that post-Brexit arrangements have created.

Regulatory Pulse

Europe and the UK are running parallel regulatory programs that address the same fraud threat through different architectures. Understanding both is essential for institutions operating across the boundary

Mandatory reimbursement in effect.

Within the regime's first 15 months, PSPs had already returned 89% of money, approx £243 million, to victims⁶⁴, converting what was once a discretionary response into a direct, predictable institutional liability that scales with payment volumes. The economic incentive to invest in pre-authorization fraud detection has never been stronger for UK FIs.

PSD3 and PSR approaching entry into force.

The EU's PSD3⁶⁵ and the accompanying Payment Services Regulation (PSR) are expected to be fully enforceable by 2027. They introduce APP fraud liability requirements, mandatory IBAN-name verification across EU payment systems, and common transaction monitoring standards across all member states. Unlike PSD2, which allowed national differences in implementation, the PSR applies uniformly across the EU, reducing the regulatory inconsistencies that fraudsters have historically exploited.

In force since January 2025	The Digital Operational Resilience Act ⁶⁶ established uniform ICT risk management, incident reporting, resilience testing, and third-party oversight requirements. Critical ICT providers, including major cloud infrastructure providers, were designated for direct EU supervisory oversight. Penalties for serious breaches reach 10% of annual global turnover or €10 million, while senior individual liability extends to €1 million. The UK runs a parallel operational resilience framework under PS24/16, in force from January 2025, but its scope and enforcement architecture differ from DORA. This creates divergent compliance requirements for institutions operating on both sides of the channel.
High-risk AI obligations	Fraud detection, credit scoring, and identity verification systems are classified as high-risk AI under the EU AI Act's Annex III ⁶⁷ , with penalties of up to €15 million or 3% of global annual turnover for non-compliance. The original enforcement date for these obligations was August 2, 2026. As of mid-2026, the proposed Digital Omnibus would extend most of these requirements to December 2027, pending formal adoption. However, obligations for general-purpose AI (GPAI) models and related transparency requirements remain on track, with enforcement beginning August 2026.

Industry Hotspots

UK retail banks and PSPs are absorbing the full economic impact of mandatory APP reimbursement while simultaneously managing the operational load of the claims. The 50/50 liability split has changed the fraud economics for receiving institutions, who were previously passive in the fraud response chain. This has created a shared incentive for cross-institutional fraud prevention that did not exist earlier.	EU payment institutions and fintechs face a dual compliance build : DORA operational resilience requirements (already mandatory) and PSD3/PSR conduct obligations (arriving 2027). Institutions that have already invested in DORA compliance have a head start on the infrastructure that PSD3/PSR transaction monitoring and fraud reporting requirements will demand. But the two frameworks require different organizational capabilities.	CROSS-BORDER MARKETPLACES Cross-border eCommerce and marketplace platforms operating across the UK-EU boundary face the most complex fraud compliance environment of any sector . This complexity is due to diverging consumer protection obligations, different data-sharing protocols, and fraud liability frameworks that are not yet harmonized across the boundary.
---	---	--

BUREAU CLOSE

Europe's regulatory stack is the most comprehensive fraud framework in the world; and it is still primarily a reimbursement framework. **It compensates victims after fraud succeeds.** The missing capability is pre-authorization intelligence: stopping fraud before the payment leaves.

That **requires continuous identity assurance, cross-institutional signal sharing, and real-time decisioning across the full customer lifecycle**. Bureau's unified identity and fraud intelligence layer is built for exactly this. It connects onboarding, behavioral monitoring, and transaction risk into a single signal that acts before authorization, not after the loss.

Southeast Asia

- **SEA Scam Compounds Export Global Operations:** Transnational syndicates⁶⁸ run AI-driven pig-butcher scams targeting global victims from border compounds.
- **Crypto Rails Accelerate Cross-Border Laundering:** Criminal networks move funds through low-regulation virtual asset exchanges to fragment transaction trails.
- **Unified ASEAN Collaboration Strengthens Ecosystem Security:** Harmonizing regulations across member states can create seamless regional oversight to reinforce systemic defenses against organized syndicates.
- **Singapore Mandates Curb Loss Volumes:** Shared responsibility frameworks dropped national scam losses by 19%, demonstrating the power of network-wide rules.
- **Philippine AFASA Standardizes OTP Phase-Outs:** Regulations mandate anti-fraud system upgrades and phase out SMS and email OTPs for high-risk transactions.

Regional Threat Signature

ORGANIZED AT NATION-STATE SCALE		RESILIENT BY DESIGN
The fraud is organized at nation-state scale. Scam compounds operate across border regions and special economic zones in Myanmar, Cambodia, and Laos. Many are sustained by local protection networks, politically connected actors, or armed groups that provide the infrastructure and security these operations require. Workers from more than 50 countries have been trafficked into these compounds⁶⁹.	The target is global, not local. Pig-butcher investment scams, romance fraud, and crypto investment schemes originating from SEA compounds target victims across North America, Europe, Australia, and the Middle East. The U.S. Department of Justice's Scam Center Strike Force⁷⁰ froze and seized more than \$580 million in cryptocurrency linked to Southeast Asia-based fraud networks in February 2026, underscoring the scale and sophistication of these operations.	The model is resilient by design. INTERPOL's Operation First Light 2026⁷¹, which ran across 97 countries and blocked 31,014 bank accounts, confirmed both the scale and the geographic reach of SEA fraud networks. It also confirmed their resilience: crackdowns displace operations into new jurisdictions, they do not eliminate the underlying infrastructure.



Top Attack Vectors

Pig-butcher investment fraud

Pig-butcher investment fraud remains the region's highest-value export. These scams rely on long-term social engineering. Fraudsters spend weeks or months building trust before directing victims to fake investment platforms. By the time money is transferred, victims believe they are investing through a legitimate platform with someone they trust. AI-personalized messaging has compressed the trust-building timeline while increasing conversion rates. The UNODC documented a 600% surge in AI-generated deepfake content linked to fraud in the region⁷², with adoption continuing to accelerate.

Scam-as-a-service infrastructure

Scam-as-a-service infrastructure supports every other attack vector. Scam centers provide more than fraud execution. They also supply the infrastructure needed to operate at scale, including script libraries, AI tools, call center operations, mule account networks, and cryptocurrency laundering channels. The ecosystem usually comprises interconnected networks of money launderers, traffickers, data brokers, and technology providers operating as a self-sustaining criminal economy. Device-level data reflects the operational footprint of this infrastructure: Bureau's fraud intelligence network found fraud exposure skewed disproportionately toward a small set of device brands relative to their overall traffic share, consistent with scam operations running standardized device setups at scale rather than organic, individually sourced hardware.

Cross-border crypto laundering

Cross-border cryptocurrency laundering takes advantage of differences in VASP (Virtual Asset Service Provider) regulatory frameworks and serves as the primary extraction mechanism. Proceeds move through cryptocurrency exchanges, operating across jurisdictions with minimal AML controls, in turn fragmenting the trail across chains and borders faster than any single regulator can follow. The INTERPOL GFFTA March 2026⁷³ assessment identifies this cross-border crypto movement as the defining financial infrastructure of the SEA fraud ecosystem.

ATO and social engineering

Account takeover and social engineering targeting domestic financial services are growing in parallel with the export model. In the Philippines, BSP data⁷⁴ shows social engineering, account takeover, and identity theft collectively accounted for 76% of total fraud losses in 2025, a profile shift that reflects both the sophistication of attack methods and the rapid expansion of digital financial inclusion.

Country Spotlights




Singapore

Singapore was the only country in the region to report a decline in scam losses in 2025⁷⁵. Losses fell 19% to S\$913.1 million; a direct result of the network-wide restrictions, rather than any single control. The Monetary Authority of Singapore's Shared Responsibility Framework and the Anti-Scam Command's **real-time fund recovery operations represent the most operationally advanced anti-fraud infrastructure in Southeast Asia**. The median loss per case rose to S\$1,644, confirming that total loss reduction came from case volume, not individual attack size.



The Philippines

The Philippines enacted the Anti-Financial Account Scamming Act (AFASA) in 2024. It is **one of the region's most substantive legislative responses to fraud**. The BSP issued implementing circulars in June 2025, setting a June 25, 2026 deadline⁷⁶ for all supervised financial institutions to upgrade fraud management systems and requiring phase-out of SMS and email OTPs for high-risk transactions by June 30, 2026. The Philippines occupies a dual position in the regional fraud ecosystem. It is both a target of domestic digital financial fraud and an operational hub linked to Southeast Asia's scam center networks.



Indonesia

Indonesia is the region's largest digital payments market by population and transaction volume, with GoPay, OVO, and Dana collectively serving millions of users. **Indonesia's financial regulator, Otoritas Jasa Keuangan (OJK), is strengthening its digital financial services framework**. However, fintech adoption is growing faster than fraud controls. This has made Indonesia an increasingly attractive target for domestic payment fraud and a transit point for mule account operations.



Vietnam

Vietnam is experiencing rapid growth in eCommerce and digital payments alongside an emerging fraud threat profile: phishing, social engineering, and investment scams are the dominant typologies. The State Bank of Vietnam has **issued fraud prevention guidance but enforcement infrastructure remains nascent** relative to the scale of the digital economy.



Thailand

Thailand faces growing cross-border exposure from scam center operations along the Myanmar border, alongside a rise in domestic online investment fraud. The Digital Economy and Society Ministry and the Royal Thai Police's Technology Crime Suppression Division are coordinating responses, but **Thailand's geographic proximity to the Mekong compound zones creates structural vulnerability** that domestic controls alone cannot address.

Regulatory Pulse

ASEAN does not have a unified fraud response framework. **Each member state has its own regulatory approach**, creating differences that transnational fraud networks are able to exploit.

Singapore's MAS framework is the regional benchmark: proactive, prescriptive, and backed by enforcement mechanisms that other ASEAN regulators are working toward. The Philippines' AFASA represents the most significant new legislative response in the region, but implementation timelines extend into 2026 and the compliance ecosystem is still being built. Indonesia's OJK, Thailand's BOT, and Vietnam's SBV are each advancing their frameworks, but from different starting points and at different speeds.

The difference in enforcement maturity between Singapore and the rest of the region is itself a vulnerability. Criminal networks route operations to the path of least regulatory resistance; and within ASEAN, that path remains clearly defined.

The most significant cross-border response mechanism in the region is INTERPOL's I-GRIP payment intervention tool, which enables real-time cross-border fund freezing. Its use in Operation First Light 2026 to block a \$6.6 million transfer in real time demonstrated that coordinated intervention is possible. But I-GRIP's **effectiveness depends on the speed and willingness of participating institutions to act**, capabilities that remain unevenly distributed across the region.

BUREAU CLOSE

The SEA fraud threat is not solvable at the platform level. Fraud originating from organized compound operations, moving through cross-border crypto rails, and targeting victims across multiple jurisdictions requires a response that no single institution can build alone.

What is needed is consortium intelligence: shared device signals, behavioral patterns, and network flows across institutions and borders that identify scam center activity before the money moves.

Bureau's cross-border identity and fraud intelligence network is designed for this environment. By connecting signals across institutions and jurisdictions, it identifies patterns that remain invisible when each organization relies only on its own data.

Asia-Pacific

- **Transnational Syndicates Exploit Jurisdictional Lacunae:** Fraud networks operate from low-enforcement bases, targeting fragmented Asia-Pacific regulatory frameworks.
- **APAC Regional Losses Reach \$9.4 Billion:** Unprecedented digital payment adoption across India, Japan, Australia, and South Korea outpaces legacy fraud controls.
- **Fraud Migrates to High-Value Digital Lending:** India sees fraud shift from UPI payment rails into digital lending as instant payment controls mature.
- **Australia Mandates Shared Cross-Sector Liability:** The Scams Prevention Framework Act enforces civil penalties across banks, telcos, and tech platforms for scam prevention failures.
- **Mandated Technical Controls Stop Account Takeovers:** Japan cut unauthorized securities access by 97% following mandatory authentication rules across brokerages.

Regional Threat Signature

Scale is accelerating across every market. Combined fraud losses across India, Japan, Australia, and South Korea reached an estimated \$9.4 billion in 2025, per India’s RBI, Japan’s National Police Agency, Australia’s National Anti-Scam Centre, and South Korea’s Financial Supervisory Service.

GEOGRAPHY	REPORTING BODY	PERIOD	REPORTED LOSS
India	RBI		₹48,021 CRORE (~\$5.6B)
Japan	National Police Agency		¥324.11 BILLION (~\$2.12B)
Australia	National Anti - Scam Centre		AUD\$2.18 BILLION (~\$1.4B)
South Korea	Financial Supervisory Service		433.8 BILLION WON (~\$300M)
Combined Losses			~\$9.4 BILLION

Each market faces a different fraud profile and each also operates under its own regulatory framework, with no unified regional coordination mechanism.

Digital financial inclusion is the primary attack surface. The rapid expansion of mobile payments, digital lending, and e-wallet infrastructure across the region has brought millions of first-time financial services users into the digital economy, and simultaneously into the fraud economy. The rapid expansion of financial inclusion has created new fraud risks that existing controls were not designed to address.

Organized transnational networks are the common thread. Whether it is pig-butcher operations targeting Indian UPI users from SEA compounds, voice phishing rings operating from Cambodia against South Korean victims, or social media investment scams targeting Japanese retirees, the operational infrastructure behind APAC fraud is overwhelmingly transnational, exploiting the absence of coordinated regional enforcement.

Top Attack Vectors

Social engineering at scale

It is the defining attack type across the region. Phone impersonation, messaging app fraud, and social media investment scams collectively account for the majority of losses. The unifying mechanic: building enough trust to move money before detection is possible.

Investment fraud via social media

It has surged across Japan and South Korea in particular with AI-generated celebrity endorsements, fake trading platforms, and cryptocurrency investment schemes targeting middle-aged savers with accumulated capital. The combination of high average loss per case and low reporting rates makes this the most economically damaging category in the region on a per-victim basis.

Digital lending and application fraud

It dominates India's institutional loss profile, as payment-level defenses have matured, fraud has migrated up the value chain into lending, where controls are comparatively thinner⁷⁷.

Mule account networks

It link all four markets with Korea and Japan both documenting transnational mule operations running from Southeast Asia, and India's RBI deploying AI infrastructure specifically to detect and disrupt mule account activity at scale.

Country Spotlights

India

India's UPI processes more real-time transactions than any other payment system in the world, and **that scale makes it as attractive to fraud actors** as it is to legitimate users. The RBI's Annual Report 2025-26 data tells a migration story. Sophisticated actors are moving away from high-volume, low-value UPI attacks and toward high-value lending fraud, evidence that payment-level detection is working well enough to push fraud elsewhere rather than eliminate it.

The scale of the RBI's response⁷⁹ reflects the scale of the problem: MuleHunter.AI now flags roughly 20,000 suspected mule accounts every month. The Indian Digital Payment Intelligence Corporation, incorporated in October 2025, is building the cross-institutional signal infrastructure that individual bank detection cannot provide.

Bureau's fraud intelligence network found **quick commerce carrying a phishing rate roughly four times that of eCommerce**, consistent with fraudsters migrating toward India's fastest-growing, lowest-friction digital commerce channels as payment-level UPI controls have matured.

The outstanding structural challenge is the NBFC and digital lending segment, where lighter-touch KYC and rapid credit extension create the same synthetic identity vulnerability that traditional banks have spent years building controls around.

Australia

Australia recorded AUD \$2.18 billion in scam losses in 2025⁸⁰, up 7.8%, with investment scams alone accounting for AUD \$837.7 million. The headline figure, however, masks a more complex picture: losses have fallen significantly from a peak of AUD \$3.1 billion in 2022. While this reflects a genuine impact from coordinated government and industry response, losses remain substantially above pre-2020 levels.

Australia's model of mandatory cross-sector obligations and meaningful penalties is attracting attention from regulators in Singapore, Japan, and South Korea as they strengthen their own scam liability frameworks.

The Scams Prevention Framework Act 2025⁸¹, which commenced February 21, 2025, is the most architecturally significant development. It **establishes mandatory obligations for banks, telecommunications providers, and digital platforms, with civil penalties of up to AUD \$50 million per violation**. The framework assigns shared responsibility for scam prevention across the three sectors most frequently targeted by fraud. Full obligations take effect March 31, 2027, with three regulators, namely: ACCC, ASIC, and ACMA, sharing enforcement responsibility.

Japan

Japan recorded its highest-ever fraud losses in 2025: ¥324.11 billion (\$2.12 billion), up 62.8% from the previous year. For the first time, the National Police Agency reported⁸² combined losses across special fraud, social media investment fraud, and romance scams.

Special fraud, mainly involving impersonation of police, prosecutors, and bank officials, nearly doubled to ¥142.3 billion across 27,832 cases, with fake police scams accounting for around 70% of losses.

Social media investment fraud reached ¥127.4 billion across 9,538 cases. About 40% of losses involved cryptocurrency, with an average loss of ¥12.1 million per case. Most victims were aged 40–60.

Japan's response shows the impact of decisive technical controls. After brokerages introduced mandatory advanced authentication, unauthorized access to securities accounts fell 97% in early 2026. The lesson is clear: **mandated technical controls can deliver measurable results within months.**

South Korea

South Korea's fraud profile is defined by voice phishing, a threat that is accelerating under AI-enabled impersonation⁸³. Voice phishing-linked bank account freezes exceeded 40,000 in 2025, with over 10,000 accounts frozen in Q1 alone.

The institutional response remains limited. Between January 2025 and August 2025, the country's five largest banks compensated only 10% of customer-reported voice phishing cases, according to FSS data⁸⁴. With government reform of the liability framework actively under discussion, the current compensation rate is likely to be a transitional figure rather than a stable policy position.

The **cross-border nature of the threat is evident**. A voice phishing network dismantled in 2026 defrauded 368 South Korean victims of ₩51.6 billion. Operating from Cambodia, Thailand, and Malaysia, it used a formal organizational structure⁸⁵ and impersonated the FSS, prosecutors, and the Korea Federation of Banks. Monthly detection of ghost accounts, exploited for voice phishing money laundering, reached 4,510 in March 2026⁸⁶, the highest since the FSS began tracking the category.

Regulatory Pulse

APAC's regulatory landscape is moving in a consistent direction: toward mandatory liability, cross-sector obligations, and AI-assisted detection; but from very different starting points and at very different speeds.

Australia is the most advanced, with the SPF Act already in force and multi-regulator enforcement operational. India is building cross-institutional AI infrastructure ahead of a formal compensation liability framework; the RBI's March 2026 consultation on unauthorized transaction compensation signals that mandatory liability is approaching. Japan's response has shifted from awareness to mandated technical controls, with the NPA and FSS both issuing sector-specific authentication requirements following the securities ATO surge. South Korea's compensation framework is under active legislative review, following evidence that bank reimbursement rates have not kept pace with reported phishing cases.

The absence of a regional coordination framework: an ASEAN or Asia-Pacific equivalent of the EU's PSD3/PSR architecture, remains the region's most significant structural challenge. Transnational fraud networks are specifically designed to exploit jurisdictional fragmentation, and no single national response can address an adversary that relocates operations across borders faster than enforcement can follow.

BUREAU CLOSE

APAC's fraud challenge lies not in technology or investment in detection infrastructure, but in coordination. **Fraud operates across borders; detection operates within them.** The disconnect between those two realities is where losses accumulate.

Bureau's cross-border identity and fraud intelligence network connects signals across institutions and jurisdictions, enabling participating organizations to identify shared identities, devices, accounts, and fraud patterns that would remain invisible in isolated data sets.

Middle East and North Africa

INTERPOL⁸⁷ rates the regional fraud risk as moderate, but that matters less than the direction it's headed. Scam operators are already relocating capacity here from Southeast Asia and they're moving faster than the region's institutional defenses can keep pace.

- **UAE Mandates Total Removal of SMS/Email OTPs:** The Central Bank forced all financial entities to eliminate SMS and email OTPs, becoming the first nation to ban them globally.
- **Saudi Arabia Issues Binding Counter-Fraud Standards:** SAMA published its Counter-Fraud Fundamental Requirements, setting mandatory control standards across banks, payment providers, and insurers.
- **Executive Impersonation Targets Regional Wealth:** AI-powered voice cloning and deepfakes target GCC corporate treasuries, sovereign wealth funds, and family offices.
- **Rapid Fintech Adoption Outpaces Fraud Controls:** Accelerating digital financial adoption across the GCC exposes users equally across low-end and premium device tiers.
- **Collaborative Data Sharing Enhances Network-Level Visibility:** Establishing regional frameworks and open data channels equips institutions with comprehensive insights to optimize proactive fraud defenses.

Regional Threat Signature

Rapid fintech adoption is outpacing fraud infrastructure.

Vision 2030 in Saudi Arabia, UAE's ambition to become a global financial technology hub, and aggressive digital payments programs across Bahrain, Kuwait, Qatar, and Oman are bringing millions of new digital financial services users online at speed. The same compressed adoption curve that created fraud exposure in India and Southeast Asia a decade ago is now playing out across the GCC, with the **added complexity of a large expatriate population, high-value target demographics, and limited cross-border enforcement coordination** within the region.

This pattern is not confined to entry-level devices or first-time users. Bureau's fraud intelligence network found high-risk session rates holding consistent at approximately 1% across low-, mid-, and premium-priced devices alike, underscoring that fraud exposure in fast-growing digital markets is not limited to any single customer segment or price tier.

Investment fraud and social engineering dominate the threat.

The GCC's concentration of high-net-worth individuals, active crypto adoption, and heavy social media use make it a high-value target for investment scams, romance fraud, and CEO impersonation attacks. AI-generated content, including fake celebrity endorsements and synthetic voice calls, is being deployed against GCC targets in the same way it is deployed against victims in Japan and South Korea.

The region's fraud reporting gap is itself a risk.

The absence of consistent, publicly reported fraud loss data across MENA creates a visibility problem for institutions making risk investment decisions. There is no equivalent of the UK Finance Annual Fraud Report or the FBI IC3 Annual Report to anchor the institutional understanding of what the fraud environment actually looks like.

Top Attack Vectors

CEO and executive impersonation fraud

Targeting corporate treasury functions is the highest-value attack vector in GCC markets. It exploits the region's concentration of large family offices, sovereign wealth vehicles, and multinational regional headquarters. AI-powered voice cloning and deepfake video have made executive impersonation operationally viable at scale. The same typology documented by Europol in European CEO fraud cases is now being deployed against GCC targets⁸⁸.

Crypto investment fraud

Is growing rapidly, targeting a population with high crypto adoption rates and a regulatory framework for digital assets that is still being built. The combination of sophisticated targets, high average transaction values, and irreversible crypto settlement creates loss profiles that are disproportionately large on a per-case basis.

Identity fraud and synthetic identity

Are emerging threats as digital onboarding expands across neobanks, BNPL providers, and digital lending platforms, particularly in the UAE and Saudi Arabia, where new financial licenses are being issued at pace. The same thin-file KYC vulnerability documented in India and Southeast Asia applies to any market where digital account opening is outpacing identity verification infrastructure.

Remittance fraud

targeting the region's large expatriate workforce, who collectively move billions of dollars monthly through formal and informal channels, represents a persistent and structurally under-addressed attack surface. Bureau's own research among senior risk and fraud leaders across the region's banking, fintech, and adjacent sectors reinforces this picture. Social engineering and phishing, alongside payment fraud, were the threats most frequently cited for 2025, often flagged by the same institutions, pointing to social engineering as a common precursor to payment loss rather than a standalone risk.

Country Spotlights

UAE

FIRST IN THE WORLD

The UAE has made **the single most decisive technical fraud prevention intervention of any country** covered in this report: the complete elimination of SMS and email one-time passwords from the financial system. CBUAE Notice 2025/3057 required all financial entities, namely: banks, finance companies, exchange houses, insurers, and payment service providers, to complete the transition to stronger authentication by March 31, 2026, making the UAE the first country in the world to formally mandate the removal of SMS OTP from financial services.

ENFORCEMENT

On the regulatory and enforcement side, Federal Decree-Law No. 6 of 2025 materially **increased the sanctions available to the CBUAE for compliance failures**. This is a structural shift from the prior framework, while extending the CBUAE's powers of early intervention, restructuring, and depositor protection. The CBUAE's 2025 Annual Report⁸⁹ frames the direction explicitly: enforcement intensity and supervisory scrutiny both escalated in the year.

STANDING

The EU removed the UAE from its own list of high-risk countries on July 9, 2025, following the UAE's removal from the FATF grey list in February 2024, both **reflecting the maturation of the UAE's AML and financial crime framework**⁹⁰. The UAE's trajectory is from reactive to architecturally proactive, building fraud prevention into the authentication and regulatory infrastructure rather than layering controls on top of existing systems.

Saudi Arabia

SAMA · APR 2026

Saudi Arabia's fraud prevention posture is anchored in Vision 2030 and SAMA's active regulatory program. SAMA published its Counter-Fraud Fundamental Requirements⁹¹ on April 13, 2026. It is a binding framework **applicable across the banking sector, finance companies, payment service providers, and insurance firms**, which establishes mandatory control standards across the full fraud risk management lifecycle.

Saudi Arabia's position is clear: **the response to financial crime is strategic, system-wide, and anchored in Vision 2030's**⁹² **commitment to integrity and transparency** as the foundation of economic transformation. The country's fraud challenge is distinctive in the regional context: the Kingdom has invested heavily in cybersecurity infrastructure and ranks among the lowest in reported fraud crimes linked to technical system breaches. However, social engineering and misuse of personal data remain the primary attack vectors; and the rapid expansion of BNPL, digital lending, and open banking creates new exposure points that the existing framework is being updated to address.

Egypt

Egypt's digital payments ecosystem is expanding rapidly under the Central Bank of Egypt's **financial inclusion mandate, with mobile wallets, instant transfers, and digital lending all growing at scale**. The fraud profile reflects the market's stage of digital maturity: social engineering and phishing targeting mobile wallet users, agent-assisted identity fraud as cash-based channels move online, and online marketplace fraud alongside the growth of eCommerce. The regulatory framework is developing but remains earlier in maturity relative to GCC peers.

Morocco

Morocco represents the North African market with the most developed financial infrastructure and the highest fintech adoption rate in the Maghreb. **Cross-border payment fraud, particularly in remittance corridors to Europe, and online investment scams targeting the Moroccan diaspora are the primary documented threat vectors**. Bank Al-Maghrib's digital payment regulatory framework is developing, but cross-border enforcement coordination with European counterparts remains limited.

Regulatory Pulse

The MENA regulatory response to fraud is accelerating from multiple directions simultaneously; and **the pace is being set by the UAE and Saudi Arabia.**

The CBUAE's OTP elimination mandate is the most technically significant regulatory intervention in the region's history. SAMA's Counter-Fraud Fundamental Requirements, published April 2026, **bring Saudi financial institutions under a binding, prescriptive control framework that aligns with international best practice.** Both represent a shift from principles-based guidance to mandatory, enforceable standards.

The region's structural challenge is cross-border coordination. MENA has no equivalent of INTERPOL's I-GRIP real-time payment intervention tool deployed at regional scale, no unified fraud intelligence sharing framework equivalent to CIFAS in the UK, and no regional equivalent of the EU's PSD3/PSR framework to harmonize conduct obligations across borders. Individual national responses, however sophisticated at the country level, cannot address a threat that is explicitly transnational in design.

BUREAU CLOSE

MENA is at an inflection point: **the digital financial infrastructure being built today will define the region's fraud exposure for the next decade.** Institutions that build fraud and identity intelligence into the foundation, rather than retrofitting it after the exposure has accumulated, have a window that more mature markets have already closed.

The problem, however, is visibility. **Without the cross-institutional data sharing that makes fraud patterns detectable at the network level, individual institution detection is limited.** Bureau's network intelligence layer provides exactly this, connecting identity signals, device intelligence, and behavioral patterns across institutions to identify fraud that no single participant can see alone.

SECTION 04

The Fraud-Resilient Playbook

CH 15


AI on Both Sides
of the Battlefield

CH 16

Model Governance: Arms Race,
Explainability and Regulatory Risk

AI on Both Sides of the Battlefield

- **Real-Time Rails Demand Sub-100ms Decisions:** Post-payment detection fails on instant rails, requiring sub-100-millisecond risk decisions to block transfers before authorization completes.
- **Multi-Modal Signal Fusion Blocks Manipulation:** Combining device, behavioral, network, and transaction data prevents attackers from bypassing controls via single-point fabrications.
- **Generative AI Industrializes Attack Creation:** Document synthesis, deepfakes, and autonomous agents allow low-skill actors to scale complex fraud using cheap compute.
- **Adaptive Adversaries Probe Model Boundaries:** Fraud networks deploy low-value attacks to map static AI thresholds, enabling automated systems to iterate around fixed controls.
- **Continuous Retraining Prevents Evasion:** Maintaining effective defense requires real-time feedback loops and ongoing model updates to counter self-adapting attack campaigns.



How AI Lowered the Barrier to Fraud

Three capabilities have made sophisticated attacks accessible to a much broader range of fraudsters.

Document synthesis


AI-generated identity documents now defeat optical inspection and automated checks at scale, a shift FATF attributes to how accessible the tooling has become⁹³. It is no longer a specialist criminal work, but a prompt and a widely available generative model.

Deepfake KYC bypass and voice cloning

Real-time face swap injection into live verification sessions and voice cloning from seconds of audio have undermined the liveness and biometric layers institutions specifically invested in to address document fraud. Europol documents both as operational tools actively deployed in CEO fraud and romance scams⁹⁴.


Autonomous fraud agent orchestration

Agentic AI systems are capable of planning, executing, and adapting multi-step fraud campaigns without human direction at each step. They represent a qualitative shift from tool to operator. The FATF has identified this explicitly⁹⁵: AI agents orchestrating fraud without human intervention. As a result, the main constraint becomes compute; and compute is abundant and inexpensive. The FBI's 2025 numbers⁹⁶ back this up with 22,364 complaints, \$893 million in losses: both likely undercounts because most victims do not recognize AI involvement.



The Defensive Architecture That Matters

Effective AI-based fraud detection follows a consistent logic regardless of the toolset:



Every layer matters. Strong models cannot compensate for weak signals, and without continuous learning, they quickly fall behind evolving fraud. Speed is non-negotiable. Fraud detection that runs after a payment is completed is already too late. Institutions operating on real-time payment rails need risk decisions in under 100 milliseconds, before authorization is granted. Any gap between payment speed and detection speed creates an opportunity that fraudsters will exploit.



Multi-Modal Signal Fusion

No single signal is sufficient. A device signal catches device anomalies but not behavioral manipulation. A transaction signal catches payment anomalies but not identity fabrication upstream. A biometric signal catches presentation attacks but not session hijacking downstream. Multi-modal signal fusion, which combines device intelligence, behavioral biometrics, network telemetry, identity graph signals,

and transaction features into a single, continuously updated risk assessment, matches the breadth of the current threat. It is also the architecture that is the hardest to attack. Simultaneously fabricating a convincing document, a consistent behavioral session, a matching device fingerprint, and a fitting network context, at real-time transaction speed, is significantly harder than fabricating any one of them in isolation.

ONE CONTINUOUSLY UPDATED RISK ASSESSMENT

Fusion matches the breadth of the current threat.

The Bank of England's April 2025 Financial Stability in Focus paper⁹⁷ framed the challenge directly: generative AI raises the ability of fraudsters to manipulate employees and customers across channels simultaneously, and institutional defense must match that multi-channel threat with multi-signal detection.

BUREAU'S FRAUD INTELLIGENCE NETWORK RECORDED

~3 fold reduction

in the share of quick commerce users on outdated Android devices between Q2 2025 and Q2 2026, an improvement not mirrored in eCommerce or matrimony platforms, where comparable rates held largely flat. This confirms device posture and session-level risk behave as distinct signals requiring separate defenses, not a single combined proxy.

The Arms Race Does Not Pause

Sophisticated fraud actors are not passively absorbing AI-based detection improvements. They are probing detection systems with low-value attacks specifically designed to map model decision boundaries. This allows them to learn what triggers a flag and calibrate future attacks to stay below the threshold. The FATF's December 2025 Horizon Scan warns that **AI-powered attack tools can adapt to defensive measures on their own**, making fraud campaigns more resilient and harder to stop.

Static AI-based detection, deployed once, managed as a fixed control, will be systematically identified and evaded over time. The adversary ecosystem, operating without regulatory constraint and with a commercial incentive to continuously improve, is built for iteration. Fraud defense is never finished. Institutions that treat AI model deployment as the end of the journey, rather than the start of continuous improvement, will fall behind adaptive fraud.

BUREAU CLOSE

AI has changed how fast the gap between winning and losing opens. **Institutions that deploy adaptive, multi-signal, real-time detection and continuously retrain against evolving attack patterns will widen their advantage.** Those that deploy AI once and manage it as a static control will find the adversary has mapped their model before the next review cycle. The battlefield is not static. Neither can the defense be.

Model Governance: Arms Race, Explainability and Regulatory Risk

- **Adversarial ML Exploits Model Blind Spots:** Fraudsters deploy low-value probes, data poisoning, and model inversion to map decision boundaries and evade detection.
- **Model Explainability Strengthens Defense:** Decision audit trails and fairness monitoring reveal probing attacks and spurious correlations before adversaries exploit weaknesses.
- **EU AI Mandates Enforce Severe Penalties:** High-risk designations require strict explainability and bias testing, backed by fines up to €15 million or 3% of global revenue.
- **Human-in-the-Loop Protocols Limit Exposure:** Formal escalation workflows for high-risk flags satisfy regulatory oversight and eliminate systemic false positives.
- **Proactive Model Governance Drives Resilience:** Integrating adversarial testing and operational transparency into core architecture beats impending regulatory deadlines.

How Models Are Being Attacked

Three adversarial techniques define the current threat to fraud detection infrastructure.

Model probing and evasion	Attackers systematically vary transaction parameters to map model decision boundaries. Once mapped, attacks are calibrated to exploit blind spots. A probed model is a model whose weaknesses are known to the adversary in advance.
Data poisoning	At institutions where fraud labels are generated through customer dispute resolution, organized actors can influence training data by systematically disputing legitimate fraud flags. If enough false negatives enter the training set, the model learns to treat attack patterns as legitimate behavior. The feedback layer, which is designed to improve the model, becomes the attack surface.
TECHNIQUE 03 Model inversion	By repeatedly querying a model and observing its outputs, sophisticated actors can reverse-engineer the features and weightings driving its decisions. A model whose architecture is reconstructible from the outside is significantly easier to evade.

The FATF observes that AI systems are now capable of autonomous adaptation in response to defensive countermeasures⁹⁸, a direct reference to the adversarial ML dynamic now operating in the fraud context.

What Regulators Are Now Asking

Regulators now expect more than accurate fraud models. They also expect institutions to demonstrate how those models are governed, monitored, tested for bias, and explained.

The EU AI Act⁹⁹ classifies fraud detection, credit scoring, and identity verification as high-risk AI systems. **That means stricter documentation, monitoring, bias testing, and explainability, not just accuracy.** Governance gaps carry a defined price: up to €15 million or 3% of global turnover. These obligations, originally due to take effect on August 2, 2026, could be delayed to December 2027 under the proposed Digital Omnibus package.

The Bank of England's April 2025 Financial Stability in Focus paper explicitly raises the need for meaningful human oversight of AI-driven financial decisions. It frames explainability as a governance prerequisite, with implications for every institution using AI in fraud detection. The FCA is moving in the same direction, with expectations hardening ahead of formal guidance.

Fraud teams operating their models as proprietary black boxes, optimized for detection performance with no audit trail of individual decisions, carry **a compliance exposure that will grow as AI regulation matures** across every major operating jurisdiction.

What Explainability Actually Requires

Explainability means that when a consequential automated decision is made, such as a transaction declined, an account restricted, or a credit application flagged, the institution can produce a defensible account of what drove it, grounded in the model's actual behavior. Three capabilities help make it possible.

02

Decision-level audit trails

Every automated decision should leave a complete audit trail, showing what signals were used, why the decision was made, and how it can be reviewed or improved.

Fairness monitoring

Fraud models should be regularly tested for unfair outcomes across different demographic groups. Historical data can introduce unintended bias, making ongoing monitoring and mitigation a regulatory expectation in the EU and UK.

Human review protocols

Not every decision should be fully automated. High-risk cases require clear escalation paths and human oversight to meet regulatory expectations and reduce the risk of unjustified harm.

The Case for Building Ahead of the Mandate

The compliance argument for building explainability infrastructure ahead of regulatory enforcement is obvious. The fraud detection argument is equally strong.

A model team that can explain what drove a decision can also identify when the model is deciding for the wrong reasons, such as:

When a spurious correlation has been encoded as a predictive feature

When a demographic variable is acting as a proxy for fraud signal

When a probing attack has corrupted the training data

The discipline that produces an explainable model also produces a more robust one.

Institutions that build explainability into the model development process will create AI systems that are easier to govern, easier to improve, and easier to defend with regulators. Those that treat explainability as a compliance exercise may satisfy an audit but gain little operational value.

The time to build this capability is now. Institutions that delay will face tighter deadlines, higher implementation costs, and less time to test and refine their AI governance.

BUREAU CLOSE

Model governance is a fraud detection function that regulators have started to scrutinize; and the scrutiny will only increase. Teams that build adversarial testing, explainability, and fairness monitoring into model development as quality standards, not compliance tasks, will be better prepared for both evolving fraud and regulatory scrutiny.

SECTION 04

The Fraud-Resilient Playbook

CH 17

The Investment Thesis:
From Cost Center to
Competitive Advantage

CH 18

Ten Strategic Priorities for
2026-2027

CH 19

The Fraud-Resilient
Enterprise Architecture
Blueprint

The Investment Thesis: From Cost Center to Competitive Advantage

- **Board Cases Demand Five-Factor Economic Frameworks:** Quantifying fraud spend requires modeling total loss exposure, growth at risk, competitive differentiation, regulatory liability, and 24-month ROI.
- **Value-Created Metrics Replace Defensive Accounting:** Shifting performance metrics from simple losses avoided to net business value created repositions fraud prevention as a strategic asset.
- **Fraud Capabilities Drive Revenue Growth:** Advanced fraud intelligence enables safer market expansion, faster customer onboarding, and higher transaction approval rates.
- **Friction Shifts Costs to Lost Revenue:** Overly aggressive fraud controls transfer financial impact from direct fraud losses to lost sales and degraded customer lifetime value.
- **Early Product Integration Reduces Long-Term Spend:** Involving fraud teams during initial product design prevents costly retrofits and minimizes customer friction at launch.



Fraud is often underfunded because organizations view it as a cost center rather than a business capability.

When fraud teams are seen as a function that consumes budget and adds operational overhead, the conversation focuses on minimizing spend: How much is needed to limit losses? What is the minimum required to meet regulatory expectations? How can costs be controlled?

These are incorrect questions. Because **they treat fraud prevention as an expense to manage rather than an infrastructure investment needed to protect and enable the business.**

Organizations that have changed this mindset are not simply spending more on fraud—they are investing differently. They recognize that strong fraud and identity intelligence capabilities do more than reduce losses. They enable safer growth, faster decisions, and greater trust.

The measurable differences between institutions at the leading edge of fraud management and those operating at the median are not primarily about technology. They are about how fraud performance is measured, where fraud sits in the organizational hierarchy, and what the function is asked to deliver.

They measure growth alongside fraud prevention

A fraud program that reduces losses by creating excessive friction is not truly effective; it simply shifts the cost from fraud losses to lost revenue and customer experience. Leading institutions track fraud rates, false positives, approval rates, and revenue impact together to understand

They treat fraud infrastructure as a growth enabler



Entering a new market, launching a new product, or serving a new customer segment requires the ability to assess risk quickly and accurately. Strong fraud and identity intelligence capabilities help institutions expand faster, reduce friction, and manage risk from day one. Without this foundation, organizations either move slowly or accept higher losses while systems adapt.

They bring fraud into product decisions early

Every new product, payment flow, or onboarding journey creates new fraud risks. Leading institutions involve fraud teams during design—not after launch because adding controls later is more costly and creates more customer friction.

They measure fraud's impact on customer lifetime value

Fraud losses are only part of the picture. Legitimate customers affected by false declines or unnecessary friction can also leave. Leading institutions measure both sides: customers lost to fraud and customers lost because fraud controls made the experience harder.



Translating fraud infrastructure investment into a board-level investment decision requires speaking the language of competitive economics rather than security risk. The five-factor framework below provides the analytical structure for that conversation.

Loss exposure	Measure the full cost of fraud: direct losses, prevention spend, customer friction, brand impact, regulatory liability, and missed opportunities. Most institutions track only a fraction of this. The board conversation should start with the true cost of fraud, not just reported losses. The scale of this unmeasured exposure is significant. Bureau's fraud intelligence network found approximately 7.4 million users, around 1.6% of a 464-million-user base, exhibiting suspicious multi-device behavior consistent with promotional abuse, a cost category that rarely appears on any standard fraud loss report.
Growth at risk	Assess what growth opportunities are limited by current fraud capabilities. What revenue could come from higher approval rates? What markets or products remain constrained without stronger risk intelligence?
Competitive differentiation	Fraud performance is becoming a competitive advantage. Institutions with stronger risk capabilities can offer better customer experiences, enter markets faster, and approve more legitimate customers. The question is: what revenue value does closing the gap with leading peers create?
Regulatory liability	Mandatory reimbursement exposure, AI governance risks, and conduct risks can be quantified before enforcement occurs. These should be modeled as financial liabilities, not treated as abstract compliance concerns.
Investment scenario modeling	Evaluate fraud investments based on measurable business outcomes over 24 months: losses prevented, approval rates improved, regulatory exposure reduced, and new opportunities unlocked. The investment case is a financial model not a security expense.

Moving fraud from a cost center to a competitive advantage requires changing the way performance is measured: from losses avoided to value created.

Losses avoided is a defensive metric. It shows how much worse outcomes could have been without fraud controls. While useful, it positions fraud as a cost function rather than a capability that drives growth, trust, and market expansion.

Value created tells a broader story: how many approvals were enabled, which markets were entered faster, how much customer value was protected, and how much regulatory risk was reduced through stronger controls.

Answering these questions requires **new measurement approaches, tracking approval rates alongside fraud losses, measuring customer impact from false declines, and quantifying regulatory exposure.** The data already exists; the missing piece is connecting it to business outcomes.

The organizations that have successfully elevated fraud to a strategic function have changed the narrative. They no longer present fraud as a problem to contain, but as a capability to build, one that delivers measurable returns and competes for investment alongside other growth priorities.

BUREAU CLOSE

The case for investing in fraud infrastructure has never been stronger. **Losses are rising, regulatory liabilities are clearer, competitive advantages from better detection are more measurable,** and the business cost of weak fraud capabilities is increasingly visible.

The question now is therefore whether the case is being framed in the business terms that make the value clear.

Ten Strategic Priorities for 2026-2027

- **Unified Architecture Closes Identity-Fraud Gaps:** CPOs and CROs must merge KYC onboarding with transaction monitoring to build continuous, multi-signal risk profiles across the customer lifecycle.
- **Dedicated Mule Detection Disrupts Scam Infrastructure:** Fraud and Compliance Officers must deploy specialized network-mapping tools to target money-laundering channels rather than just payment symptoms.
- **Networked Threats Mandate Cross-Sector Consortia:** Fraud and legal executives must join cross-institutional intelligence networks to uncover shared device, identity, and money mule signals.
- **Proactive Defense Requires AI Attack Red-Teaming:** CISOs must regularly pressure-test models against deepfake video injections, adversarial model probing, and training data poisoning.
- **Board Governance Demands Full Economic Modeling:** CFOs must measure total fraud impact, including direct losses, customer friction, and regulatory liabilities, to justify strategic 24-month investment cases.

The fraud threat is moving faster than most institutions are responding to it. These ten priorities are not a comprehensive fraud program. They are the decisions that, made in the next 18 months, will determine which institutions lead that response and which continue to follow it. Each priority is sequenced by urgency. Each carries an ownership assignment because priorities without owners are observations, not decisions.

Unify Identity and Fraud Intelligence Architecture

The separation between identity verification and fraud detection is a major structural weakness. When these systems operate independently, critical signals are missed, a synthetic identity can pass KYC without risk context, while **later fraud signals may not connect back to the identity layer**. The solution is not replacing systems, but creating a shared identity intelligence layer that connects onboarding, authentication, behavioral, and transaction signals into a continuous risk view.

OWNER

Chief Product Officer + Chief Risk Officer

Implement Real-Time Fraud Decisioning Infrastructure

Institutions operating real-time payments without sub-100-millisecond fraud scoring have a critical gap. Fraud can be completed in the seconds between authorization and detection. Closing this gap requires **risk decisions at the speed of the payment rail**, before approval, not after settlement. For institutions using FedNow, Faster Payments, UPI, PayNow, or PromptPay, this is one of the highest-priority fraud prevention investments.

OWNER

Chief Technology Officer + Head of Fraud

Build a Dedicated Mule Detection Program

Mule detection is not the same as fraud detection. It requires a different approach, tracking account behavior, mapping networks, and sharing intelligence across institutions. Transaction-level fraud models alone cannot identify the infrastructure behind fraud proceeds. Institutions that **detect fraud but miss mule networks are only addressing the outcome, not the source**. The solution is a dedicated mule detection capability with specialized signals, continuous feedback, and reporting to industry fraud databases.

OWNER

Head of Fraud + Chief Compliance Officer

Harden Biometric Verification Against AI Injection Attacks

Biometric KYC and liveness systems without cryptographic device attestation remain vulnerable to face-swap injection attacks, a known attack method, not a future risk. Vendors should **demonstrate device attestation, active challenge-response, and protection against injection attacks** as core security requirements. The audit starts with one question: what prevents a third-party tool from injecting a deepfake video feed into the verification process?

OWNER

Chief Product Officer + Chief Technology Officer

Invest in Explainable AI Governance Infrastructure

The EU AI Act is an immediate driver for explainability, but it is not the only reason to build it. **Explainable models help teams identify when decisions are based on flawed signals, biased patterns, or compromised data.** Decision audit trails, fairness testing, and human review processes are the minimum requirements. They should be built as model quality standards—not just compliance documentation.

OWNER

Chief Risk Officer + Chief Technology Officer

Engage Fraud Intelligence Consortia

Synthetic identity fraud, mule networks, and cross-border fraud operations often remain invisible to individual institutions because the signals are spread across the ecosystem. Detecting them **requires a shared view of identities, behaviors, and transactions over time**, something no institution can build alone. Cross-institutional fraud intelligence networks are the structural response to this challenge. The governance, legal, and data-sharing frameworks can be solved. The alternative: fighting networked fraud with isolated data, is no longer sustainable.

OWNER

Head of Fraud + General Counsel

Model the True Cost of Fraud for the Board

Direct fraud losses rarely reflect the full cost of fraud. Prevention spend, customer friction, regulatory exposure, brand impact, and missed growth opportunities all add to the true cost, but are often missing from standard fraud reporting. Build a complete fraud cost model using the six-layer framework. The goal is not to create urgency, but to **establish the right baseline for investment decisions.** The case for fraud infrastructure changes when the board sees the full cost.

OWNER

Chief Financial Officer

Build a Regulatory Early Warning Capability

The regulatory shifts shaping fraud and AI governance have not happened without warning. PSD3/PSR, the EU AI Act, Australia's SPF, SAMA's Counter-Fraud Fundamental Requirements, and India's unauthorized transaction compensation consultation all provided months or years of advance signals through consultations, policy updates, and implementation timelines. **Institutions that engaged early built readiness ahead of enforcement;** those that waited are now building under pressure. A dedicated regulatory monitoring function that translates policy developments into operational actions is a small investment compared with the cost of reactive remediation.

OWNER

Chief Compliance Officer

Pressure-Test Fraud Controls Against AI Attacks

Adversarial ML techniques such as model probing, data poisoning, and model inversion are not future risks—they are active attack methods. Any fraud model that has not been tested against these techniques has an unknown level of exposure. Regular red-team exercises should **test against real attack methods, including AI-generated document fraud, deepfake liveness bypass, session injection, and probing attempts** designed to map model behavior. The findings may be uncomfortable, but they provide a more accurate view of risk than audits based on the assumption that existing controls are working.

OWNER
Chief Information Security Officer

Develop a Fraud Talent and Capability Strategy

The combination of fraud expertise and AI/ML capability is one of the scarcest skills in financial services. Fraud teams understand the problem but often lack AI depth; AI teams build models but often lack fraud context. The **lack of integration between the two impacts detection quality**. A strong fraud talent strategy goes beyond hiring; it builds AI skills within fraud teams, embeds fraud expertise into model development, and creates the cross-functional capability needed for adaptive detection. Institutions that keep fraud and data science separate risk building technically strong models that fail in real-world fraud environments.

OWNER
Chief Human Resources Officer + Head of Fraud

BUREAU CLOSE

As fraud is a permanent feature of any system where value can be transferred, **these ten priorities are about strengthening institutional capabilities** to keep pace with adversaries investing in more advanced attacks. The opportunity to close that divide is limited.

The Fraud-Resilient Enterprise Architecture Blueprint

- **Event-Driven Assessment Secures Transactions:** Real-time evaluation of user actions, devices, and transactions blocks fraudulent transfers before funds clear.
- **Unified Ownership Eliminates Identity-Fraud Silos:** Consolidating KYC onboarding and fraud detection under one executive creates connected risk decisioning across the customer lifecycle.
- **Identity Graphs Expose Complex Fraud Networks:** Linking device telemetry, behavioral biometrics, and transaction histories over time exposes synthetic identities and account takeovers.
- **Continuous Feedback Powers Adaptive Detection:** Instrumenting real-time feedback loops and ongoing retraining prevents model drift against evolving attack patterns.
- **Explainable Audit Trails Satisfy Regulatory Scrutiny:** Capturing exact signals, scores, and rules behind automated decisions provides the transparency needed to defend AI systems.



Fraud resilience is an architecture¹⁰⁴, not a single control. Leading organizations embed risk intelligence across the customer lifecycle: from onboarding and authentication to transactions and account management. This allows them to create a connected defense against evolving threats.

Organizational Architecture

The organizational **dimension of fraud resilience offers the greatest opportunity for improvement**, and requires cultural change as much as structural change.

Unified Leadership

Fraud and identity are the same problem approached from two directions. Identity verification asks: is this person who they claim to be? Fraud detection asks: is this person behaving consistently with who they claim to be? However, **managing these functions separately leads to fragmented decisions and creates blind spots that fraudsters exploit.**

Fraud resilience starts with unified ownership. One executive oversees identity, onboarding, and fraud across the customer lifecycle, enabling the integration that separate teams rarely achieve.

Cross-Functional Fraud Council

Every product decision is also a fraud decision. Every new payment flow, onboarding journey, or credit product creates new fraud risks, and adding controls after launch is far more costly than building them in from the start.

A cross-functional fraud council helps prevent these costs by involving product, engineering, legal, compliance, and customer experience teams early in the design process. **Fraud intelligence shapes product design, while product priorities shape fraud controls**, ensuring both evolve together, not one after the other.

Technical Architecture

The technical architecture of a fraud-resilient enterprise follows a consistent design logic: every customer interaction generates a signal, **every signal contributes to a persistent identity record**, every record feeds an adaptive model, and every model decision is explainable, auditable, and available for regulatory review.

Event-Driven Risk Assessment

Static fraud detection relies on historical data and isolated transactions. It cannot keep pace with real-time payments or account takeovers that unfold in seconds. Event-driven risk assessment evaluates every user action, device event, login, and transaction in real time, updating risk before a payment is approved.

This **shift from batch processing to event-driven detection is essential for real-time payments**. Leading institutions assess risk before money moves; others detect fraud only after it has happened.

Identity Graph Persistence

The identity graph is the foundation of continuous identity assurance. It brings together every customer signal, such as device fingerprints, behavior, authentication history, network context, and transactions, into a single, continuously updated view across the customer lifecycle.

Its purpose is to establish what normal looks like for each customer so unusual activity stands out. Synthetic identities lack a genuine history, and account takeover attackers cannot immediately replicate a customer's normal behavior. **By connecting signals over time, the identity graph exposes fraud patterns that transaction-level detection alone cannot see.**

Adaptive Model Layer

A fraud model trained once and left unchanged quickly falls behind as attack methods evolve. The adaptive model layer keeps detection current through continuous feedback from confirmed fraud and legitimate activity, performance monitoring to detect model drift, and regular retraining against new attack patterns.

Bureau's fraud intelligence network recorded device-based fraud holding steady at around 3.9% of all sessions across Q1 and Q2 2026. The consistency shows that **device-based fraud is a persistent risk rather than a seasonal spike**, reinforcing the need for continuous model retraining instead of periodic refreshes.

Adaptation requires a feedback loop that is instrumented, monitored, and acted upon: a dedicated model operations function that treats detection quality as a continuous improvement objective rather than a deployment milestone.

Explainable Decision Audit Trail

Every automated fraud decision should include a complete audit trail, capturing the signals, model scores, rules, and outcome behind the decision.

This serves three purposes: it explains decisions to customers and regulators, helps improve models by identifying decision errors, and provides the evidence needed to demonstrate that AI systems are operating as intended, fairly, and with appropriate human oversight.

How the Two Architectures Work Together

The organizational and technical foundations of fraud resilience must work together. Unified leadership without shared data creates fragmented decisions. Adaptive models without real-time threat intelligence fall behind evolving attacks. And, explainable AI without cross-functional collaboration meets compliance requirements but does little to improve product decisions.

The **blueprint works as an integrated system**. Each capability adds value on its own, but together they create a fraud program that learns from every interaction, adapts to new threats, and delivers the coordination and transparency that today's risk and regulatory environment demands.

BUREAU CLOSE

Fraud resilience is **an operating model that requires continuous investment, adaptation, and coordination** across the business. The organizations that lead in fraud prevention have built systems that learn from every interaction, adapt to new threats, and respond faster than attackers.

Bureau is built on these principles. By bringing together device intelligence, behavioral biometrics, identity graph analysis, and real-time risk decisioning on a single platform, Bureau helps institutions operationalize this architecture across the customer lifecycle.

Ready to assess your fraud resilience? Speak with Bureau to evaluate your current architecture and identify the capabilities needed for the next generation of fraud defense.

APPENDIX

Bureau's Unified Risk Decisioning Platform

Fraud risk now spans the entire customer journey, from onboarding to every transaction. A synthetic identity approved at onboarding can become a mule account weeks later, while the same device, identity, or network may surface across multiple institutions.


Most fraud stacks operate in silos, with separate tools for identity, devices, behavior, transactions, AML, and investigations. Each control sees only part of the user's journey, making it harder to connect signals, detect coordinated fraud, and act before losses occur.

Organizations need one risk layer that connects identity, device, behavior, transaction, telco, and network intelligence across the entire customer lifecycle. It should let teams build, test, deploy, and monitor risk workflows from one place, while continuously carrying forward everything the platform knows about each user.

How Bureau's Unified Risk Decisioning Platform Addresses These Challenges


Bureau brings fraud, identity, and compliance capabilities into one decisioning engine and one workspace. Teams can build workflows, test them against live-like payloads, deploy through an API, SDK, or hosted flow, and monitor outcomes without switching systems.

Every interaction: from the first onboarding check to login, device change, beneficiary addition, and payment, updates a continuous risk state. Network intelligence adds context beyond the organization's own data, while automated decisions can allow, step up, or block in under 50ms.



The Value Bureau Brings

One continuous risk state:	Identity, device, behavior, transaction, and network signals build on one another across the lifecycle.
Faster risk operations:	Build, test, launch, and tune workflows without waiting for engineering releases.
Earlier fraud detection:	Risk signals act from the first interaction, before fraudulent accounts or transactions progress.
Network-level intelligence:	Connect fraud patterns across identities, devices, accounts, and institutions.
Smarter investigations:	Link alerts to sessions, entities, and fraud rings instead of investigating events in isolation.
Complete control and visibility:	Every decision, rule change, alert, and case action is versioned, explainable, and auditable.



Level up from account-level defense to network-level intelligence. See Bureau’s Unified Risk Decisioning Platform in action.

Book a demo at marketing@bureau.id

BUREAU FRAUD INTELLIGENCE REPORT 2026

91

Account Takeover (ATO)

When a fraudster gains control of a legitimate account, often through stolen credentials or session hijacking, enabling unauthorized transactions, data theft, and further fraud.

Adversarial Machine Learning

Techniques used to probe, manipulate, or reverse-engineer fraud detection models, helping fraudsters find and exploit weaknesses in automated defenses.

AFASA (Anti-Financial Account Scamming Act)

A Philippine law requiring financial institutions to strengthen fraud controls and phase out SMS/email OTPs for high-risk transactions, raising security standards for digital payments.

AML (Anti-Money Laundering)

Laws and controls designed to prevent, detect, and report money laundering, helping stop criminals from moving or disguising illegal funds.

APP Fraud (Authorized Push Payment Fraud)

When a victim is tricked into authorizing a payment to a fraudster's account, resulting in losses even though the payment was technically authorized.

Behavioral Biometrics

Using how a person types, swipes, navigates, or interacts with a device to verify identity, helping detect suspicious users without adding friction.

BNPL (Buy Now, Pay Later)

Short-term credit that lets customers pay for purchases in installments, creating a target for fraudsters using synthetic identities.

Bust-Out (Synthetic Identity)

When a fraudster maxes out the credit available to a synthetic identity before abandoning it, causing significant losses in a short period.

CBUAE (Central Bank of the UAE)

The UAE's central bank and financial regulator, responsible for overseeing the country's banking and financial system.

Cifas

A UK fraud prevention organization that shares fraud data and operates the National Fraud Database, helping organizations identify known fraud patterns and individuals.

CNP Fraud (Card-Not-Present Fraud)

Payment fraud where the physical card is not used, such as online or phone purchases, making transactions harder to verify.

Continuous Identity Assurance

Ongoing verification based on behavior and other signals throughout the customer lifecycle, helping detect account compromise after onboarding.

Credential Stuffing

Automated attempts to use stolen username-password combinations across websites, exploiting reused credentials to take over accounts.

DORA (Digital Operational Resilience Act)

EU rules requiring financial institutions to strengthen technology risk management, incident reporting, and third-party oversight, improving resilience against technology and cyber risks.

Data Poisoning

Manipulating training data so a fraud model learns incorrect patterns, potentially causing it to overlook fraudulent activity.

Deepfake

AI-generated audio, video, or images that convincingly impersonate a real person, potentially bypassing identity checks or enabling social engineering.

Device Fingerprinting

Creating a profile from a device's technical characteristics to identify and distinguish devices, helping link devices to suspicious or past activity.

Device Spoofing

Falsifying a device's identity or characteristics, such as through emulators or virtual environments, to evade fraud and security controls.

EU AI Act

EU rules that classify certain AI uses, including some fraud and identity systems, as high-risk, requiring greater oversight, testing, documentation, and transparency.

Explainability (Model Explainability)

The ability to explain why an AI model made a particular decision, supporting regulatory compliance, audits, and accountability.

FATF (Financial Action Task Force)

The global body that sets standards for preventing money laundering and terrorist financing, shaping how countries regulate financial crime.

FCA (Financial Conduct Authority)

The UK regulator responsible for overseeing financial services and protecting financial markets and consumers.

Fraud-as-a-Service

Fraud tools, stolen data, fake identities, and other capabilities sold or rented to criminals, making fraud easier to launch and scale.

FedNow

The US Federal Reserve's real-time payment system, enabling faster movement of funds and reducing the time available to stop fraudulent payments.

Federated Learning

A method that lets organizations train a shared AI model without sharing their raw data, enabling collaboration while better protecting sensitive information.

GDPR (General Data Protection Regulation)

EU data protection law governing how personal data is collected and used, creating privacy requirements that can affect cross-border fraud intelligence sharing.

Identity Graph

A map linking identities, devices, behaviors, authentication, and transactions, revealing hidden connections and fraud patterns across the customer lifecycle.

I-GRIP (Global Rapid Intervention of Payments)

An INTERPOL mechanism for rapidly freezing funds during cross-border fraud investigations, helping stop stolen money from being moved or withdrawn.

IOCTA (Internet Organised Crime Threat Assessment)

Europol's annual assessment of organized online crime and cyber threats, providing insight into emerging criminal methods and trends.

KBA (Knowledge-Based Authentication)

Identity verification using personal information or security questions, a method increasingly outdated because much of this information is available to criminals.

KYC (Know Your Customer)

The process of verifying a customer's identity, typically during onboarding, helping prevent impersonation and other identity fraud.

Liveness Detection

A check that confirms a real person is present rather than a photo, video, or synthetic replica, helping prevent biometric spoofing.

MFA (Multi-Factor Authentication)

Authentication requiring two or more independent verification methods, providing stronger protection than passwords alone.

MFA Fatigue

Repeated authentication requests sent to a user until they approve one, allowing attackers to bypass security through social engineering.

Model Inversion

Repeatedly querying an AI model to infer the information or logic behind its decisions, helping attackers understand and potentially exploit the model.

Model Probing

Using small test transactions to learn how a fraud model makes decisions, allowing attackers to adjust future activity to avoid detection.

Money Mule

A person or account used to receive and move illegal funds, helping criminals hide the source of stolen money.

MuleHunter.AI

An RBI-deployed AI tool that detects suspected mule accounts using behavioral analysis, helping banks identify accounts used to move illicit funds.

NBFC (Non-Banking Financial Company)

A financial institution that provides services such as lending without a full banking license, playing a major role in India's financial system.

OJK (Otoritas Jasa Keuangan)

Indonesia's financial services regulator, responsible for overseeing financial institutions and markets.

Onboarding Fraud

Fraud committed when an account is created using a stolen, fabricated, or synthetic identity, allowing criminals to enter a financial or digital platform undetected.

Payment Rail

The infrastructure through which payments are processed and settled, such as UPI or FedNow, determining how money moves between accounts.

Pig-Butchering Scam

A long-term scam in which fraudsters build trust before directing victims to fake investments, often resulting in large financial losses.

Point-in-Time Verification

Checking an identity only once, usually at onboarding, leaving organizations with limited visibility into identity changes or account compromise later.

Payment Rail

The infrastructure through which payments are processed and settled, such as UPI or FedNow, determining how money moves between accounts.

Pig-Butchering Scam

A long-term scam in which fraudsters build trust before directing victims to fake investments, often resulting in large financial losses.

Point-in-Time Verification

Checking an identity only once, usually at onboarding, leaving organizations with limited visibility into identity changes or account compromise later.

Privacy-Enhancing Technologies (PETs)

Technologies that enable data to be analyzed or shared without exposing raw personal data, supporting fraud intelligence while protecting privacy.

PSD3 / PSR (Payment Services Directive 3 / Payment Services Regulation)

Forthcoming EU payment rules that include APP fraud liability and IBAN-name checks, increasing fraud prevention and reimbursement obligations for payment providers.

PSP (Payment Service Provider)

A company that provides payment services to businesses or consumers, enabling payments to be processed and received.

PSR (Payment Systems Regulator, UK)

The UK regulator overseeing payment systems and the APP fraud reimbursement regime, shaping how payment providers manage fraud liability.

RBI (Reserve Bank of India)

India's central bank and primary financial regulator, responsible for overseeing banks, payments, and financial stability.

Real-Time Payment Rail

A payment system that moves funds almost instantly, often leaving less time to detect and stop fraud.

Regulatory Arbitrage

Exploiting differences between countries' rules by routing activity through weaker regulatory environments, creating gaps fraudsters can exploit.

SAMA (Saudi Central Bank)

Saudi Arabia's central bank and financial regulator, responsible for overseeing the country's banking and financial system.

Scam-as-a-Service

Criminal services such as scripts, call centers, and money-laundering channels sold to fraudsters, making large-scale scams easier to operate.

Scam Compound

A facility where organized scam operations are run, sometimes using forced labor, enabling large-scale fraud against victims worldwide.

Session Hijacking

When an attacker steals an already-authenticated session instead of login credentials, potentially bypassing passwords and MFA.

SIM Swap

When a fraudster transfers a victim's phone number to a SIM they control, allowing interception of OTPs and account recovery codes.

SPF Act (Scams Prevention Framework Act, Australia)

Australian legislation requiring banks, telecom providers, and digital platforms to strengthen scam prevention, creating mandatory obligations and potential penalties.

Synthetic Identity

A fake identity combining real and invented information, allowing fraudsters to open accounts or obtain credit without using a single stolen identity.

UPI (Unified Payments Interface)

India's real-time payment system, enabling instant account-to-account payments at massive scale.

Voice Cloning

AI-generated speech that convincingly copies a person's voice, enabling impersonation and social engineering.

VASP (Virtual Asset Service Provider)

A business providing services such as cryptocurrency exchange or transfers, subject to varying financial crime and AML requirements across jurisdictions.

1, 5, 30, 32, 48, 50, 61 Europol, IOCTA 2026, April 2026:

25, 31, 95, 98. FATF, Horizon Scan: AI and Deepfakes, December 2025:

3, 6, 28, 39, 51, 52, 54, 56, 96. FBI IC3, 2025 Internet Crime Report:

4, 8, 73, 87. INTERPOL, GFFTA, March 2026:

7, 24, 41, 57, 93. FATF, Cyber-Enabled Fraud Paper, February 2026:

9, 16, 22. FTC, Congressional testimony, March 25, 2026:

10, 12, 58. UK Finance, Annual Fraud Report 2026:

11, 42, 44, 45, 77, 78. RBI, Annual Report 2025–26:

13, 19, 59, 64, 100. PSR, APP Scams Reimbursement Dashboard Q4 2025:

14, 36, 75. Singapore Police Force, Annual Scam and Cybercrime Brief 2025:

15. RBI Governor Sanjay Malhotra, public statement, November 2025 · Ministry of Home Affairs:

17. PSR, One Year On: Impact of APP Reimbursement on Victims, October 2025:

18, 103. DLA Piper / CMS GDPR Enforcement Tracker, 7th Edition:

20, 81, 101. Scams Prevention Framework Act 2025:

21. RBI, March 2026:

23. Federal Reserve Bank of Boston, Gen AI Is Ramping Up the Threat of Synthetic Identity Fraud, April 2025:

26, 62, 88, 94. Europol, EU-SOCTA 2025:

27, 97. Bank of England, Financial Stability in Focus: Artificial Intelligence in the Financial System, April 2025:

29, 49. FBI IC3, PSA251125, November 2025:

33, 38, 71. INTERPOL, Operation First Light 2026, July 9, 2026:

34, 43, 63. Cifas, Fraudscape 2026, June 2026:

35. FCA, Firms' Use of the National Fraud Database and Money Mule Account Detection Tools, December 3, 2025:

37, 79. RBI Governor Sanjay Malhotra, November 2025:

40, 55. Federal Reserve Bank of Richmond, Buy Now, Pay Later: Recent Developments and Implications, February 2026:

46. FTC, Data Spotlight: Social Media Scams, April 2026:

47. FTC, Consumer Sentinel Network Data Book 2024:

53. Federal Reserve Financial Services, December 2025:

60. UK Finance, Half Year Fraud Report 2025:

65. Council of the EU / European Parliament, April 23, 2026 · Morrison Foerster analysis:

66. EIOPA:

67, 99, 102. EU AI Act, Article 99:

68, 72. UNODC, Inflection Point: Global Implications of Scam Centres, Underground Banking and Illicit Online Marketplaces in Southeast Asia, April 21, 2025:

69. UNODC, Strategic Infiltration of Vulnerable Jurisdictions Through Criminal Foreign Direct Investments:

70. US DOJ, Scam Center Strike Force, February 2026:

74. BSP, as reported by Philippine Daily Inquirer, February 5, 2026:

76. BSP / CICC, June 2025 · BSP AFASA implementing circulars:

80. National Anti-Scam Centre, Targeting Scams Report 2025:

82. Japan National Police Agency, 2025 Fraud Statistics (Preliminary), February 12, 2026, as reported by The Japan Times:

83. Korea Times, September 2025, reporting FSS data:

84. Korea Herald, October 2025, reporting FSS data:

85. Seoul Economic Daily, May 2026:

86. Herald Economy / BigGo Finance, April 2026:


89. CBUAE, 2025 Annual Report, April 9, 2026:

90. Herbert Smith Freehills Kramer, September 2025:

91. SAMA, Counter-Fraud Fundamental Requirements, April 13, 2026:

92. SAMA Governor, speech, October 16, 2025:

104. Bureau, Unified Risk Decisioning Platform:



BUREAU

© [2026] Bureau id. All rights reserved.



www.bureau.id



hello@bureau.id