GROUP-IB

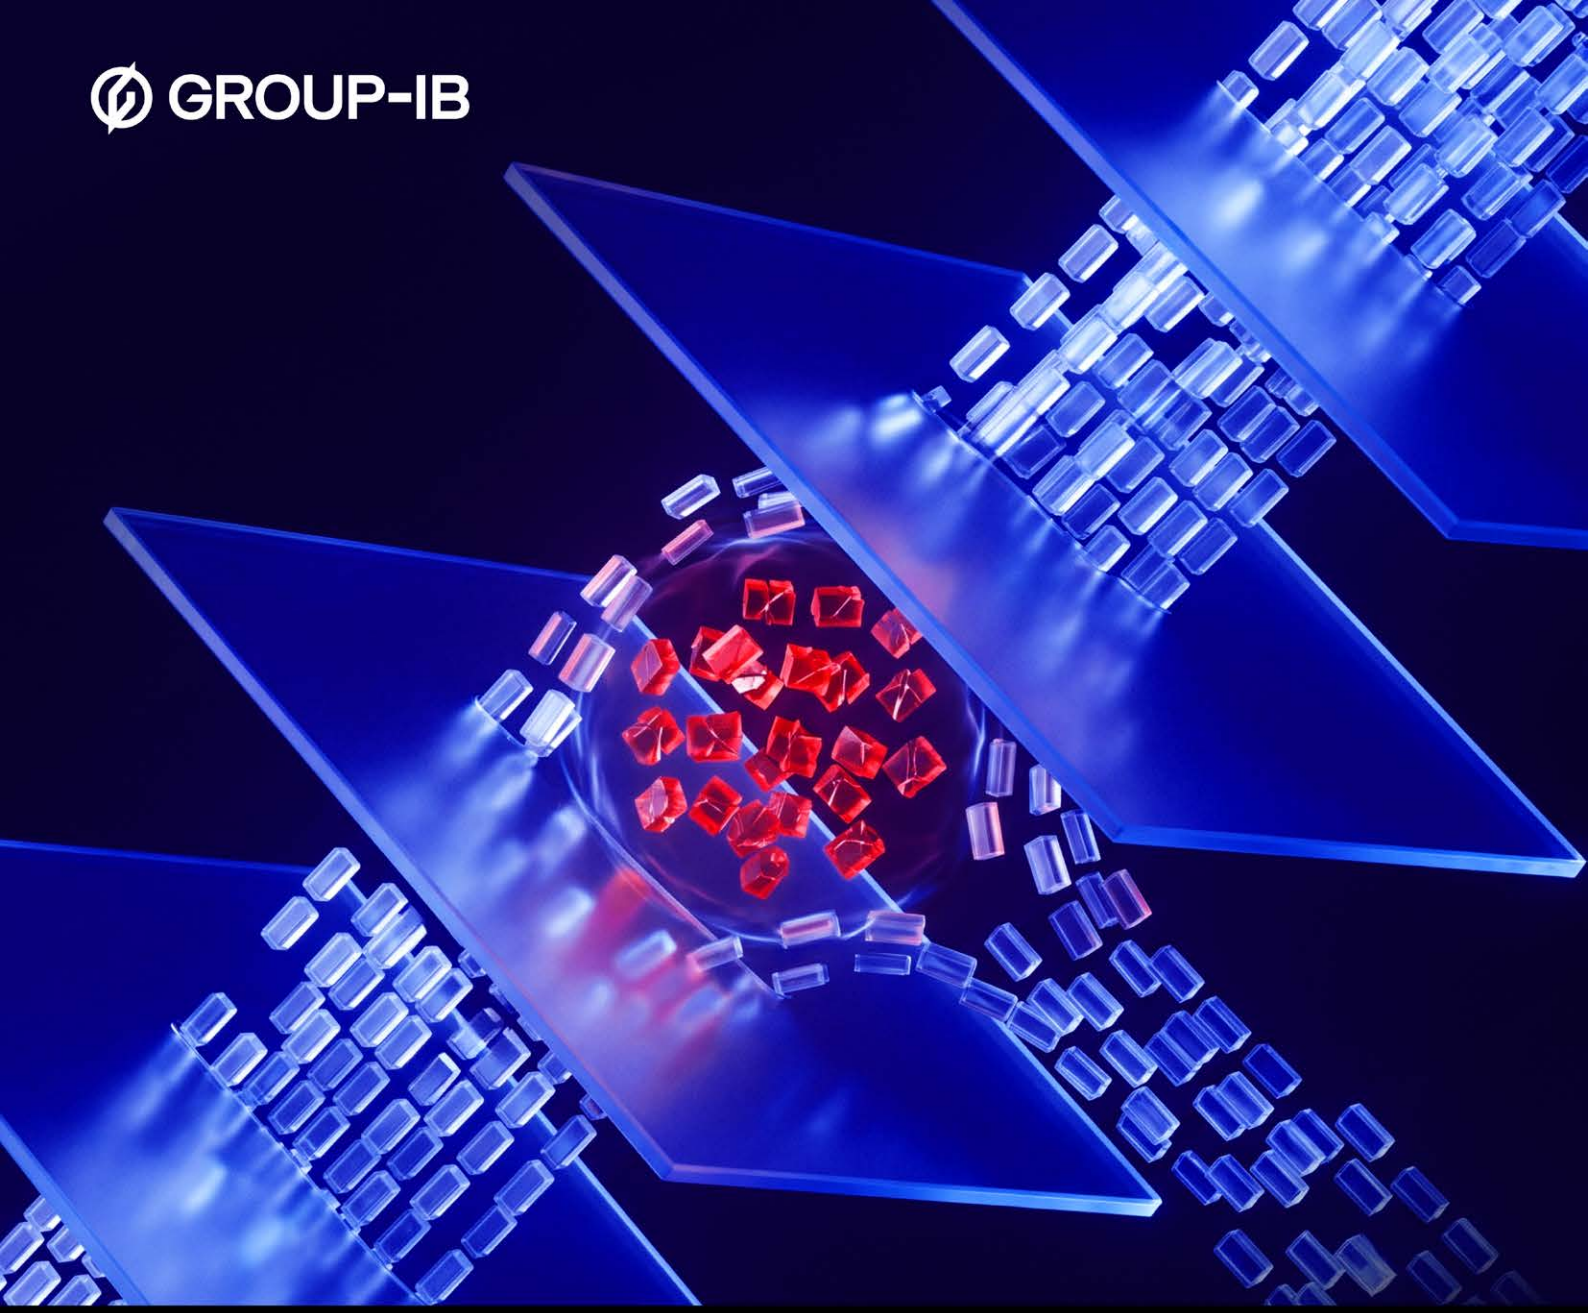

White Paper

Before Fraud Transacts

Enabling Proactive Prevention for European Finance
(In the Age of AI)

GROUP-IB.COM

Table of contents

Executive Summary

European financial institutions are navigating an unprecedented intersection of three converging forces: **the industrialisation of fraud enabled by artificial intelligence, an increasingly prescriptive regulatory environment demanding real-time monitoring and cross-institutional collaboration, and a structural intelligence gap that leaves most organisations defending against individual events while adversaries coordinate entire campaigns.**

This whitepaper synthesises original research, operational intelligence from Group-IB investigations, and emerging regulatory requirements to define what **proactive fraud prevention must look like in 2026 and beyond.**

Key Findings

40%
EU fraud attempts involve AI
Group-IB 2025 research[1]

28%
Funds exit mule accounts in 15min
RUSI 2025 data

378%
Growth in synthetic ID fraud
EU Financial Sector[2]

60%
Fraud loss reduction possible
With predictive analytics[3]

PSR/PSD3 creates a direct balance-sheet liability for PSPs that fail to implement continuous, session-wide fraud monitoring — making pricing model selection a risk management decision, not a procurement one.

Fraudsters have shifted from high-volume, low-value campaigns to precision strikes : fewer attacks, far greater financial impact per operation.

AI-driven deepfakes, synthetic identities, and dark LLMs remove the language, scale, and cultural barriers that once limited European exposure.

Over 40% of fraud attempts on EU financial institutions now involve AI; synthetic identity document fraud has grown 378% in recent cycles.

The Intelligence-Sharing Paradox prevents timely intervention: 28% of fraudulent funds leave a mule account within 15 minutes, while “confirmed fraud” data is only shareable after 24+ hours.

Cyber, Fraud, AML, three separate teams, routinely investigate the same adversary campaign in isolation, each reaching partial conclusions, leading to incomplete protection.

01

The New Economics of Fraud

The economics of fraud have fundamentally shifted. What was once a volume game (spray-and-pray campaigns optimised for reach) has transformed into a professionalized, industrialised operation optimised for impact. Attackers now invest in infrastructure, tooling, and tradecraft before targeting a single victim.

1.1 Emerging Fraud Types and Trends Across Europe

The fraud landscape entering 2026 is defined by a single, decisive shift: attackers have moved from high-volume, low-value campaigns to professionalized, high-impact operations. The era of randomized fraud is giving way to precision strikes: fewer attacks, but far more damaging ones. Several converging trends are driving this transformation across European markets.

£2,427

Avg loss per APP fraud case[4]
UK Finance Annual Fraud Report 2025

8,065

Deepfake KYC bypass attempts at a single EU institution, Jan–Aug 2025
Group-IB Weaponized AI 2026[5]

These aren't isolated attack types. They're a supply chain. A synthetic identity passes KYC today. It will become the account holder in a mule network next month.

Deepfake-as-a-Service (DaaS): High-Impact Impersonation

Deepfakes became widely available as a service in 2025, enabling even low-skilled criminals to deploy hyper-realistic voice and video cloning at scale. These are no longer niche tools, they are commoditised attack vectors available on underground markets, priced per campaign.

Deepfake investment and Authorised Push Payment (APP) scams use AI to create realistic investment opportunities, leading to a 23% jump in average loss per case.

AI-Enabled Authorised Push Payment (APP) and Social Engineering Chains

Authorised Push Payment (APP) fraud is accelerating as AI removes the human limitations of social engineering. Automated tools execute

conversation flows across thousands of potential targets simultaneously, sensing hesitation, adjusting language in real time, and applying micro-manipulations — subtle behavioural cues that suppress inhibitions and push victims toward action.

The Fraud Ecosystem Is Interconnected by Design

The categories below are not separate problems; they are nodes in the same criminal network. Most institutions still defend against individual events rather than coordinated operations.

Key Fraud Vectors

- Payment fraud (card-not-present)
- Authorization flow risks & account takeover
- Identity fraud & synthetic identity clusters
- Invoice fraud & business email compromise

Systemic Enablers

- Deepfake voice/video impersonation
- Dark LLMs & hybrid AI-human operations
- Money laundering & terrorist financing
- Insider threats & internal embezzlement

The thread connecting all vectors: adversaries chain these attack types into coordinated campaigns. Detection built around individual categories will always arrive late to an operation that has already moved on.

Why APP Fraud Works: The Exploitation of Human Decision-Making

Every fraud prevention control is designed around a system. APP fraud is designed around a person.

Victims are not careless or uninformed. The attack is engineered to defeat awareness, using fear, urgency, and perceived authority to override rational decision-making precisely when it matters most. A spoofed number. Transaction details harvested from prior reconnaissance. A manufactured emergency that makes hesitation feel dangerous. Group-IB has observed this model scale beyond account-level fraud in several European markets — criminals sustaining multi-day manipulation campaigns that pressure victims into liquidating major assets entirely.

No transaction monitoring system catches a payment that looks completely legitimate, made by the genuine account holder, on their usual device, at a normal time. APP fraud passes every signal-based control because the signals are all correct. The attack happened before the transaction.

The intervention window is earlier — in the infrastructure being staged before the call is placed, in the scam script being tested across institutions, in the phone number being conditioned for impersonation. That is where intelligence-led prevention operates. Not on the transaction. On what preceded it.

1.2 How AI Supercharges Fraud Sophistication

AI-assisted tooling, attack enablement, and technique improvisations have given adversaries leverage greater than anticipated. Previously, commoditised attack campaigns were built for scale and speed. Now, attackers prioritise personalisation and precision, where one timely, well-crafted attack can cause disproportionate impact.

When it comes to phishing and social engineering, automated AI tools will not only execute ploys at unprecedented speed and scale, enabling conversation flows across thousands of targets simultaneously, but introduce subtle nuances previously impossible.

They can sense hesitation, pauses, and behavioural cues, tailoring messages in real time to suppress inhibitions and push victims toward action.

Impersonation, Deepfakes & Synthetic Identities

Deepfake-as-a-Service growth

233%

YoY increase in unique selling usernames (2023–2024)

+52%

2025 trending

Consequences

$347 mln

in verified deepfake fraud losses in Q2 2025

KYC bypass attempts

8,065

deepfake fraud attempts (Jan–Aug 2025)

80%

occurred in final four months of period

KYC bypass tools list

- DEEPFACELIVE
- LDPLAYER
- VOLCAM
- XPOSEDBRIDGE

Synthetic identities

AI-generated CVs, profiles, voices, and interviews used to pass hiring checks

300+

companies inflitrated via synthetic remote workers

Consequences

- FRAUD
- DATA THEFT
- LONG-TERM PERSISTENCE

Source: Hi-Tech Crime trends 2026

AI Removes the European Advantage

European industries serve diverse populations, each with different languages, cultural contexts, and norms. What were once harder to target at scale, aren't anymore.

AI REMOVES THIS FRICTION ENTIRELY:

- Language barriers → gone
- Cultural nuance → exploited
- Scale without local knowledge → possible

Group-IB conducted a probing operation to test public resilience against rising deepfake and impersonation attacks. **The findings were stark — over 40% of fraud attempts on European financial institutions now involve AI, with deepfakes used to impersonate account holders and bypass authentication controls.**[6]

Synthetic Identity Fraud: Automated, Adaptive, Self-Learning

Synthetic identity fraud is not just impersonation — it is the creation of whole identities used for fraud, account takeover, and access. These AI agents behind fabricated identities change behaviour based on approval/decline outcomes, adjust timing, content, and channels, and learn from friction points such as Multi-Factor Authentication (MFA) challenges, verification steps, and transaction blocks.

While addressing concerns around synthetic media, with fraud losses climbing into the hundreds of millions of euros and synthetic identity document fraud growing by 378%, **Group-IB's vishing experiment exposed how cybercriminals exploit AI voice cloning and weaknesses in telecom systems.**

AI-generated voices can reliably mimic real speech patterns from just seconds of source audio. Dark LLMs (self-hosted AI models) and hybrid human-AI operations are already being used to craft scam scripts, assist voice scams, and automate social-engineering workflows at scale.

> From the frontlines of cybercrime, we see AI giving criminals unprecedented reach. Today it helps scale scams with ease and hyper personalisation at a level never seen before. Tomorrow, autonomous AI could carry out attacks that once required human expertise. Understanding this shift is essential to stopping the next generation of threats.

Anton Ushakov,
Head of Cybercrime Investigations Unit, Europe, Group-IB

1.3 Regulatory Pressure for Greater Granularity and Accountability

The regulatory shift underway across Europe is not incremental. PSR and PSD3 don't encourage better fraud controls, they create direct financial liability for PSPs that fail to implement them. The burden of proof has reversed: under PSD2, banks routinely denied refunds by arguing customer negligence. PSR reverses that. The bank must now prove the customer acted fraudulently. In spoofing scenarios, that bar is significantly higher than most institutions' current evidence architecture supports.

The practical consequence is immediate. Article 83 of PSR requires continuous, session-wide monitoring covering device intelligence, malware detection, remote access tool presence, behavioural biometrics, and location anomalies across every customer session, not just flagged ones.

Once enacted, selective monitoring will fall short of the compliance threshold. Institutions whose fraud controls rely on high-risk transaction triggers must fundamentally rethink their prevention model they simply haven't faced the liability event that will make that visible yet.

APP scams no longer sit outside the AML risk framework either. As PSR and PSD3 take effect and AMLA shapes supervisory standards, institutions that treat APP fraud as a fraud team problem rather than an AML exposure will face consequences on both fronts simultaneously.

1.4 The Unsustainable Cost Curve of Fraud Prevention

The November 2025 provisional agreement on PSR and PSD3 marks a decisive shift: PSPs that fail to implement adequate fraud controls are now directly liable for customer losses. Fraud prevention is no longer a cost centre — it is a balance-sheet risk.

The API Pricing Trap

Article 83 sets an outcome-based mandate: every PSP must operate transaction monitoring mechanisms that prevent and detect fraudulent payment transactions, analysing the environmental and behavioural characteristics typical of each customer and factoring in, at a minimum, compromised credentials, known fraud scenarios, and signs of malware in authentication sessions. PSPs whose fraud prevention mechanisms prove inadequate are expected to cover the resulting customer losses.

Meeting that outcome across the fraud tactics driving today's losses, such as remote access scams, social engineering, and account takeover, is difficult without broad signal collection: device intelligence, behavioural analysis, and remote access detection across the customer journey. This is where the economics turn against the PSP. Under event-based pricing models, every additional signal is an additional API call. The vendor bill grows in direct proportion to monitoring thoroughness.

The result is a perverse incentive: a fraud system doing its job well costs more money. PSPs caught in this model face a choice between rationing signal collection to control costs, accepting the fraud losses and liability exposure that creates, or monitoring comprehensively and absorbing an uncapped, unpredictable cost curve.

Compliance Thoroughness vs. Cost: The Misalignment

Under event-based pricing, the cost of fraud monitoring scales with the very behaviour PSR Article 83 demands. The deeper you monitor, the more events you generate, and the more you pay. A flat-rate model breaks that link.

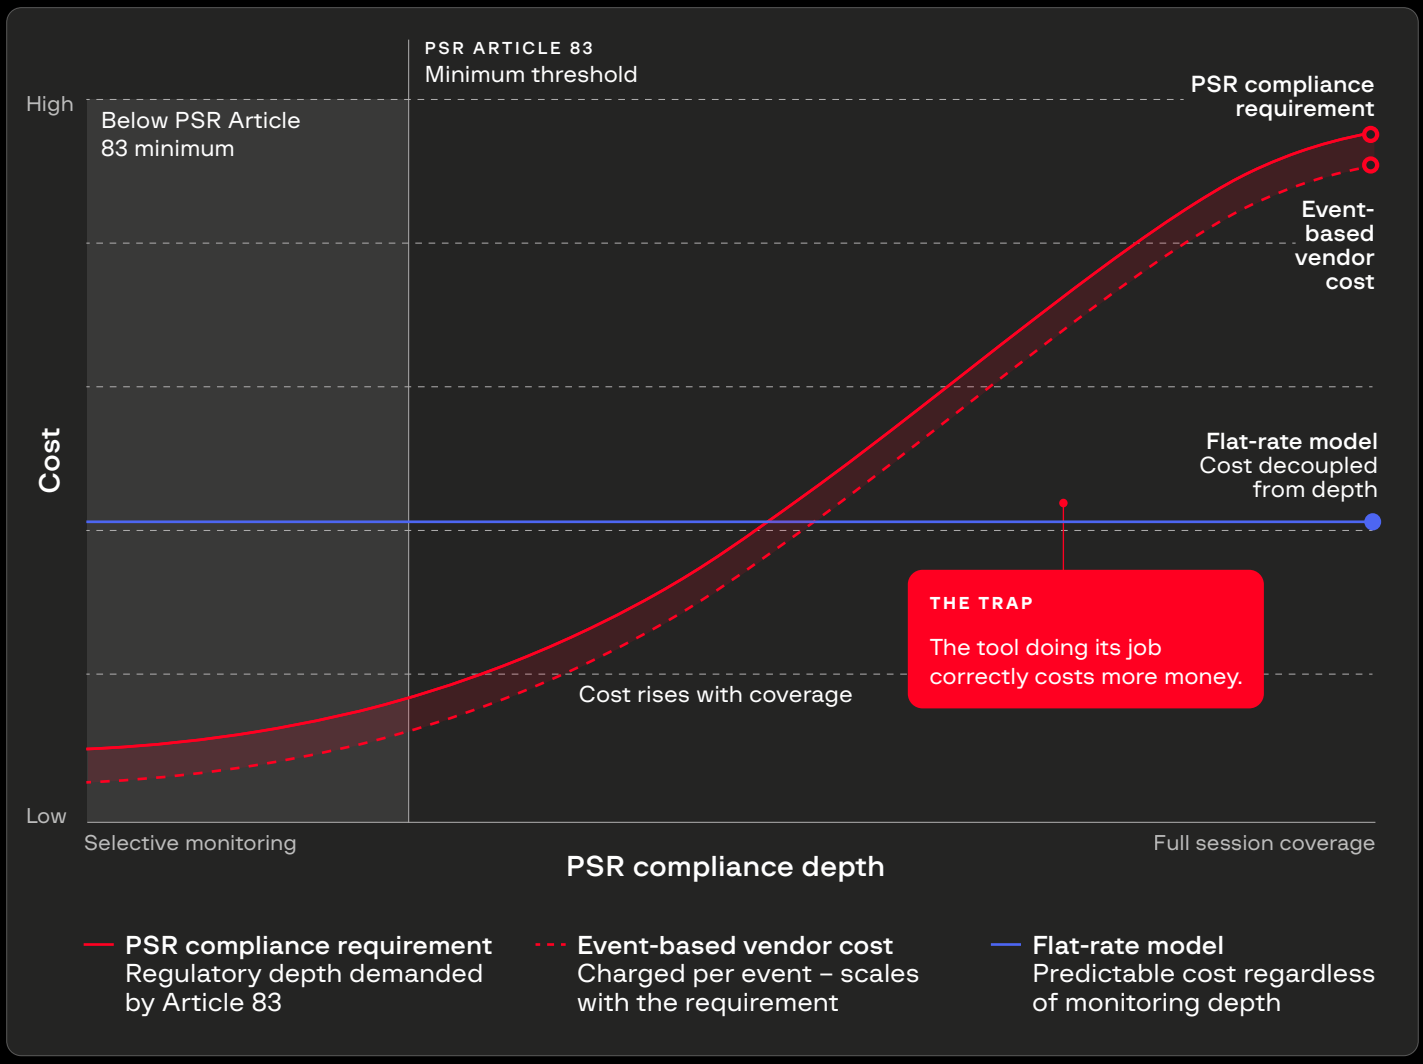

THE LOCK

01

Cost and compliance move together. **The two red lines are virtually identical** because event-based pricing scales with the same activity PSR demands.

THE PENALTY

02

Every step toward fuller session coverage **widens the bill**, not the risk gap. Compliance becomes the cost driver.

THE RELEASE

03

A flat-rate model **breaks the relationship** – depth of monitoring no longer determines spend.

False Positives and Customer Friction

False positive rates compound this. Without rich contextual signals, systems work with a fraction of the available picture.

43%

Early rule-based fraud detection rate

(With 15%+ false positive rate)[7]

39.6%

Accuracy improvement from device fingerprinting

(Across mobile and web)[8]

60%

Reduction in fraud losses with rich signals

(While cutting false positives 50%)[9]

PSPs are caught between two expensive outcomes: pay more for context and control false positives, or ration context and pay the cost in customer attrition and compliance exposure. Under PSR, that contextual blind spot carries a regulatory price tag that did not exist under PSD2.

1.5 The Geopolitical Dimension: State-Aligned Fraud Targeting European Finance

Most European fraud teams model their adversary as a financially motivated criminal network. **Group-IB's 2026 research requires updating that assumption.**

The boundary between financially motivated cybercrime and state-directed operations effectively disappeared in 2025. North Korean-aligned groups ran large-scale cryptocurrency theft against financial infrastructure — at times layering ransomware as a secondary revenue stream. Group-IB tracked 626 active corporate access listings targeting European organisations on dark web markets last year, priced and available to any downstream actor regardless of motivation.[10]

The most documented convergence: **DPRK IT workers, threat actors tracked by Group-IB**[11] under aliases including Wagemole and Jasper Sleet, using fabricated identities, deepfaked video interviews, and AI-generated employment histories to infiltrate financial institutions as legitimate employees. Primary objective: salary generation for the North Korean state. Secondary objective: system access, data familiarity, and facilitation of broader DPRK cyber operations. Over 300 companies confirmed to have infiltrated in 2025.

For European fraud leaders, the operational consequence is specific: threat models that treat fraud, insider risk, and Advanced Persistent Threat (APT) activity as separate disciplines will miss operations designed to look like one while being another. The answer isn't separate teams with separate tools. It's unified signal visibility across cyber, fraud, identity, and payments — applied before the operation completes.

02

Blind Spots in Today's Fraud Intelligence

Inside most financial institutions, the battle against financial crime is fought by three teams that rarely share the same intelligence picture. Fraud prevention, AML compliance, and cybersecurity each maintain their own systems, alert queues, data sets, and definitions of what constitutes a risk event. The criminal on the other side of the screen has no such organisational boundaries.

50%

Organisations with fully siloed Fraud & AML tools

Live polling 2025, The Paypers[12]

55%

Companies working in silos

Industry research 2025[13]

54%

Financial leaders: fragmentation blocks crime detection

Tookitaki research[14]

2.1 **Fragmented Intelligence:** Data Silos Across AML, KYC, and Fraud Monitoring

The mechanics of a typical attack illustrate exactly how this plays out.: An Account Takeover (ATO) occurs, followed by layering fraud — multiple transfers, pseudo and mule accounts, cross-border transactions. Which team owns the detection and mitigation? Is it cyber, because the attack started with a breach? Fraud, because suspicious logins and money movement were involved? Or AML?

The answer is all of them, but each team investigates a different fragment of the same crime. If cyber stops the login, fraud blocks the payment, and AML files the report, each system registers an alert within its own environment. Without entity resolution and cross-domain correlation, no system recognises the full adversary topology. This is a fragmented response against a staged, coordinated attack chain.

Today, cyber, fraud, KYC, and AML teams hold related indicators but cannot correlate them. Different teams often reach different conclusions about what is risky based on limited visibility.

The impact of this fragmentation is measurable. Once a scam is activated, only 20% of APP fraud cases are detected and reported within the first two hours[15]. Even within a 24-hour window, reporting does not reach 60%. By then, funds have moved, infrastructure has shifted, and recovery becomes exponentially harder.

Why the Silos Persist

AI-related fraud techniques are becoming more convoluted, with adversaries now orchestrating multiple trails, attack vectors, and stages within a single campaign. Indicators are stored, processed, and scored in separate systems, without a shared analytic model connecting identity, device, transaction, and infrastructure signals:

- Cyber scores intrusion risk
- Fraud scores transaction anomalies
- AML scores structuring patterns

Each unit initiates its own response to a shared problem. The infrastructure persists. The adversary campaign remains operational. The defence, as a result, is reactive, prolonged, and often non-strategic.

Against adversaries pivoting in real time and purposely introducing smokescreens for strategic diversions, fragmented visibility is not just inefficient. It is a structural advantage handed to the attacker.

The Path to Convergence

Group-IB Fraud Matrix addresses this directly – its Fraud Intelligence mapping the full fraud kill chain to identify and disrupt the entire fraud ecosystem (from the actors, infrastructure, tools, and distribution channels that enable fraud) before the fraud succeeds. It provides operational intelligence highlighting the most-used tactics and the stages where teams must converge their activities.

First Order of Business: Data Convergence

- AML transaction monitoring
- KYC identity intelligence & risk signals
- Device & behavioural intelligence
- Cyber and fraud telemetry
- Unified into a single internal dashboard

Second Order: Closing the Feedback Loop

- Cross-team alert correlation
- Shared entity resolution
- Continuous controls upgrade
- Adversary campaign tracking
- Evolving against fraudster tactics in real time

Industry Spotlight: F3 — Fighting Fraud Through Collaboration

As the fraud industry moves toward greater standardization, Group-IB is helping shape the conversation. The company serves as an initial data contributor to the newly released MITRE Fight Fraud Framework™ (F3) developed by the MITRE Corporation. Much like ATT&CK transformed cybersecurity operations by establishing a common language for adversary behavior, F3 aims to provide a standardized taxonomy for fraud techniques. Group-IB's contribution reflects its extensive experience investigating fraud operations globally and reinforces its commitment to advancing collaborative, intelligence-led fraud defense across the financial ecosystem.

2.2 Limited Integration With Cybersecurity Teams

This moves beyond a data-sharing problem to decision fragmentation. Limited integration within teams often manifests in separate triggers, escalations, and entity resolution because the organisation responds to alerts, not to adversary campaigns.

The mechanics illustrate this precisely: a cybersecurity team receives intelligence about rising malware infections on consumer devices. Because the threat does not directly affect the bank's own infrastructure, no alert is passed to the fraud team. Weeks later, the fraud team notices a surge in suspicious payment activity. Without context from the cybersecurity team, they treat it as isolated fraud. By the time the root cause, which is device-level malware enabling session takeover, is identified, the fraud has already been monetised and the funds moved.

Showcasing the delay in detection sequence vs fraud movement

DAY 1

A threat intelligence analyst detects a new mobile banking trojan circulating via fraudulent app store listings. Because no bank-owned infrastructure is directly compromised, no cross-team alert is raised to the fraud function.

DAY 12

The malware, now active on an estimated 3,200 consumer devices, begins harvesting session tokens and intercepting One-Time Password codes during live banking sessions. Fraud monitoring sees no anomalous transaction signal because no transaction has yet occurred.

DAY 23

Fraud analytics flags a surge. 280 accounts showing high-velocity transfers to first-time payees within a 36-hour window. The fraud team investigates the transactions in isolation, attributing the pattern to a social engineering campaign.

DAY 31

CERT shares an indicator of compromise matching the malware variant. Cyber and fraud teams, working from separate case queues, independently close their respective investigations — cyber marks it "contained," fraud marks it "resolved with partial recovery.

TOTAL CONFIRMED LOSS

€1.8M. Time from first malware signal to fraud team awareness: 23 days. Cross-team investigation sessions held: zero. This is not an edge case. It is the architectural default in most institutions today.

2.3 Insufficient Industry- and Brand-Wide Intelligence Sharing

Even when institutions overcome internal silos, a second, overarching challenge comes to light: the non-dissemination of intelligence across organisations — **what we term the external cohesion gap.**

Threat actors do not create phishing kits, hosting infrastructure, synthetic identities, or mule accounts to target only one enterprise. Their targets are distributed, and attack manoeuvres are pivoted to the next victim if the first attempt is unsuccessful. Infrastructure (domains, mule networks, synthetic identities, automation frameworks) often persists and is tested across multiple institutions.

Individually, what appears to be a successful mitigation may, in the broader ecosystem, strengthen adversary infrastructure because no one collectively viewed the network topology, cross-bank routing patterns, shared device fingerprints, reused onboarding anomalies, or infrastructure clustering.

If we are fighting 15-minute crimes with 24-hour data, we aren't stopping fraud. We are documenting losses.

The Intelligence-Sharing Paradox

According to recent RUSI research, 28% of fraudulent funds exit a mule account within just 15 minutes. Over half are gone within an hour. Yet the industry standard is to share "Confirmed Fraud" lists. Confirming fraud and filing a SAR typically takes 24 hours or more.

Sharing intelligence about suspicious activity before fraud is confirmed has always meant unacceptable compliance risks. Traditional methods like SHA-256 hashing create re-identification vulnerabilities that violate privacy laws. And by the time fraud is confirmed and shareable, the money is gone.

Consider Authorized Push Payment (APP) fraud. When customers themselves authorise transactions, fraud becomes much harder to detect. Even with the right device, at the right time, and in the right location, decisions can be manipulated. APP fraud is no longer just about system vulnerabilities; it's about social engineering and exploiting trust.

APP Fraud

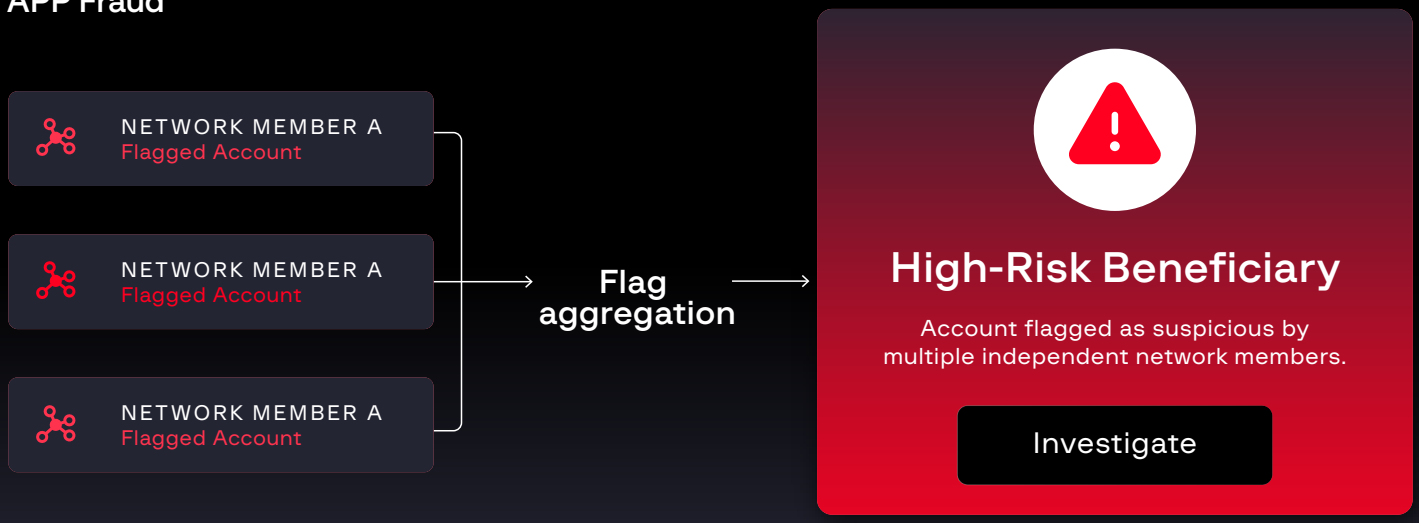

Detected Aggregation Points

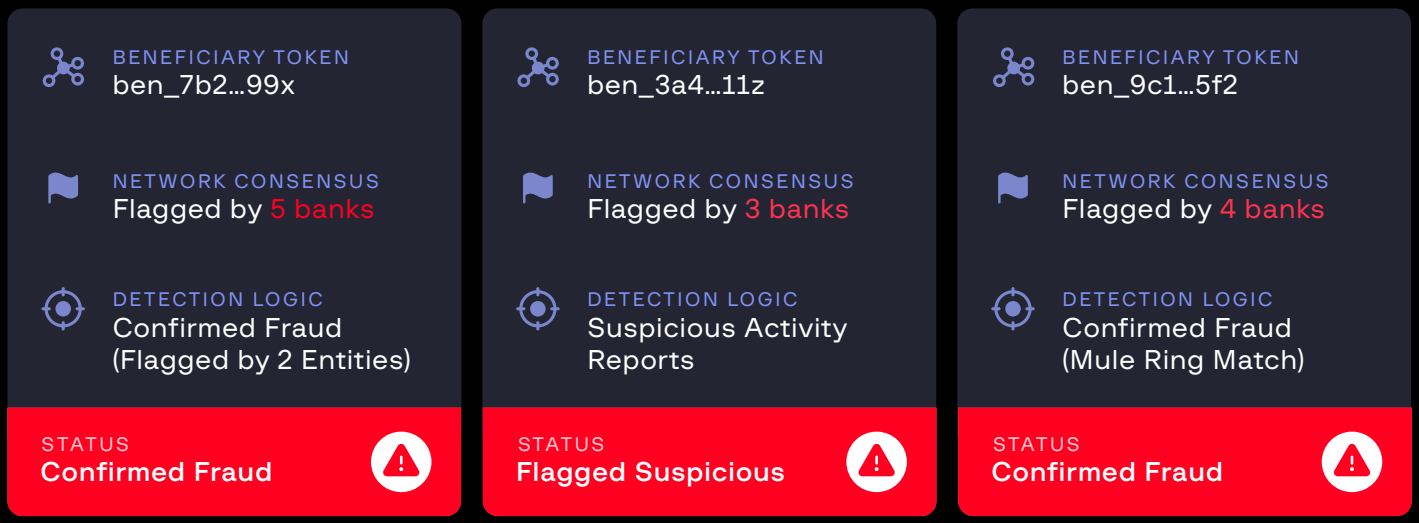

Image: Cross-institutional flag aggregation identifying a High-Risk Beneficiary through network consensus across multiple independent member banks.

By the time suspicious activity clears the threshold for confirmed fraud and becomes shareable, the transaction has long been authorised, the funds have moved, and the window to intervene has closed. This is the paradox made tangible. What is needed is actionable, cross-institutional intelligence that is anonymised, infrastructure-focused, privacy-preserving, and real-time.

At Group-IB, we describe this as the Intelligence-Sharing Paradox: regulatory frameworks encourage collaboration against financial crime, while privacy and data protection requirements limit how intelligence can be shared in practice.

Resolving the Paradox:
Group-IB Cyber Fraud Intelligence Platform

What is needed is actionable, cross-institutional intelligence that is anonymised, infrastructure-focused, privacy-preserving, and real-time. Group-IB's Cyber Fraud Intelligence Platform (CFIP) uses patented Distributed Tokenisation, validated by **Bureau Veritas** as GDPR-compliant to eliminate re-identification risks while enabling institutions to share suspicious signals during the critical warm-up period.

For the First Time, Institutions Can:

- Share suspicious signals during the critical warm-up period — before fraud is confirmed
- Detect coordinated attacks across multiple institutions in real time
- Stop APP scams, mule networks, and synthetic identity fraud before funds are transferred
- Meet regulatory collaboration requirements (UK PSR, Singapore COSMIC, EU PSD3) without compliance exposure

2.4
The Gap Between Cyber Intelligence and Payment Decisions

What Cyber Intelligence Sees

- Known fraud infrastructure
- Mule account networks
- Scam campaign activity
- Phishing kit deployments
- Synthetic identity clusters
- Ransomware wallet connections
- Dark web marketplace links

PAYMENT APPROVED HERE
without the intelligence

What Payment Teams See

- Beneficiary account number
- Transaction amount
- Time and location
- Device signal
- Sender behaviour

The question is no longer "does this transaction look suspicious?"
It is "does this recipient sit inside an active criminal ecosystem before funds clear?"

Group-IB Suspicious Payment Module
Connects transaction data with external threat intelligence at the moment of payment decision – so every approval is made with full ecosystem visability

The Ecosystem Risk Chain

Modern laundering isn't random. A single approved payment connects outward through an entire criminal infrastructure.

01
Approved Payment
Looks routine. Right device, right location, right time. No flag raised.

02
Scam Campaign
Beneficiary account tied to an active social engineering or investment scam operation.

03
Phishing Kit
Credentials harvested through reused phishing infrastructure across multiple institutions.

04
Synthetic Identity Cluster
Fabricated or stolen identities used across onboarding, layering, and account control.

05
Mule Account Network
Funds move through layered mule accounts across jurisdictions. 28% exit within 15 minutes.

06
Ransomware Wallet
Value converted and routed into crypto infrastructure linked to ransomware operations.

07
Dark Web Marketplace
Credentials, access, and stolen data listed and resold. The cycle begins again.

WITHOUT ECOSYSTEM VISIBILITY

Node 01 is all you see.

Traditional transaction monitoring flags anomalies in isolation. By the time the chain is traced, funds have moved through five layers and three jurisdictions.

WITH GROUP-IB INTELLIGENCE

The chain is visible at Node 01.

The suspicious payment module connects beneficiary accounts to known fraud infrastructure before approval — so the full ecosystem risk is assessed at transaction time.

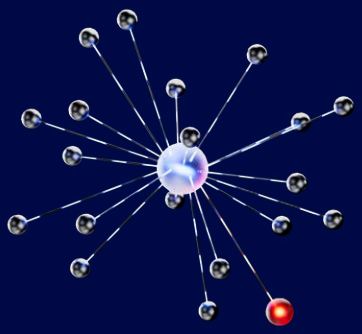

Image: A showcase of what is missed as a consequence of the decision gap

The divide between cyber intelligence and payment decisions is costing the industry billions. Threat intelligence teams see the infrastructure. Payment teams see the transaction. Rarely do the two views converge in time.

Traditional transaction monitoring was not designed to detect ecosystem-linked risk in real time. This leads teams to make decisions without full visibility on money flow. As a result, fraudulent payments get approved. The challenge is no longer spotting a “suspicious payment,” it is understanding whether the recipient sits inside an active criminal ecosystem before funds clear.

A New Risk Layer in Modernised Infrastructure

This problem compounds as European banks modernise payment rails and tokenise assets. As more institutions experiment with tokenised deposits, stablecoin custody, and crypto trading, new vectors emerge like tokenised asset fraud, stablecoin-based layering, and cross-border crypto settlement abuse. Value moves faster, and laundering becomes more automated. Detection windows shrink accordingly.

Group-IB’s Suspicious Payment Details module in its proprietary **Threat Intelligence** addresses this directly. Rather than only flagging suspicious transactions in isolation, it flags accounts associated with known fraud infrastructure, identifies mule accounts linked to scam ecosystems, and connects payment beneficiaries to underground criminal networks. So, a routine payment can be assessed against the full criminal context surrounding the recipient.

2.5 The Rise of Internal and Insider-Enabled Fraud

Modern banks handle massive volumes of digital transactions and uphold privileged access for many employees, vendors, and contractors, creating multiple potential vectors for insider abuse. Insider fraud is among the top types of financial crime affecting institutions, alongside external frauds like APP and identity theft.

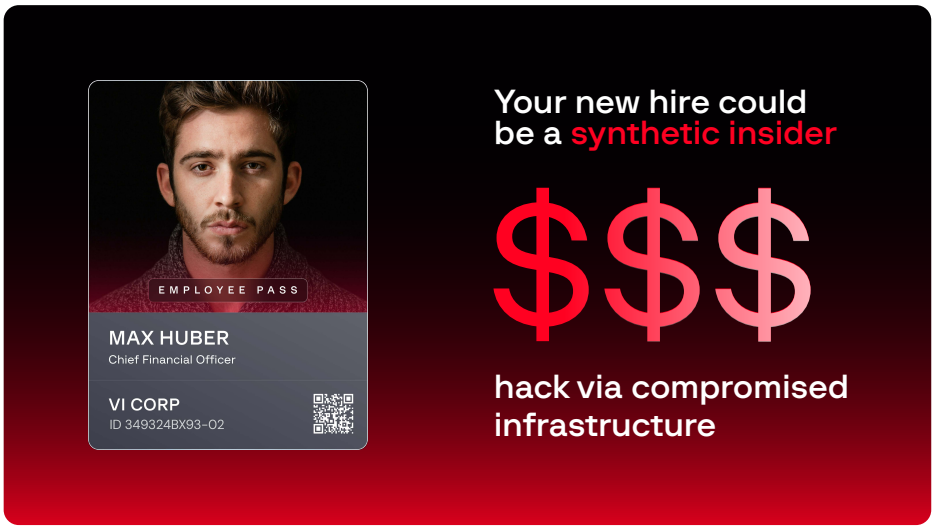

The Shift:
Why Break In When You Can Be Let In?

Legitimate-looking identity abuse is emerging in two major ways. **The first is structural: implanted through fictitious, fabricated employees, accelerated by GenAI and deepfake identities are inserted into vendor ecosystems to exploit trust.**

Attackers are not breaking in; they are onboarded. They move through trusted integrations with legitimate permissions within your system.

The second is operational: employment fraud, where a legitimate employee outsources work and creates proxy access from undisclosed locations or through VPN masking.

Attack Pattern Shift

EARLIER PATH
External actor → exploit → breach → privilege escalation

NOW
Fabricated identity → onboarding → legitimate access → ecosystem spread

The perimeter may hold. Identity does not.

Identity at the Intersection of Cyber and Fraud

Modern attacks increasingly leverage identity abuse — whether through manipulation, malicious insiders, synthetic employees, or proxy access — rather than purely external attack vectors. The question is no longer "How do attackers break in?" but "How does legitimate access become weaponised?"

Detection built around external threat models — perimeter breaches, anomalous login attempts, known malware signatures — will not catch insider-enabled fraud. What is required is continuous, cross-domain visibility that treats identity behaviour and payment activity as a unified signal, monitored from onboarding through offboarding.

Taxonomy of an insider fraud:
A 90-day synthetic employee compromise

What can a synthetic insider disrupt within your organization in less than a quarter? Let's find out.

Organisation: 40,000+ employees | Hybrid workforce | Multiple SaaS integrations | Vendor-connected supply chain

Day 0–14

PHASE 01

Entry Through Trust

A contractor is hired for a remote engineering role.

- ✓ Valid ID documents submitted
- ✓ Clean background check
- ✓ Strong technical interview performance
- ✓ Equipment shipped to the registered address

No perimeter breach. No phishing email. No exploit. The individual onboarded is not the person working. Access granted:

- Corporate email + VPN
- Git repository
- Internal collaboration tools
- Limited SaaS access

From a security perspective, Everything appears compliant.

Day 15-30

PHASE 02

Access Stabilisation

The synthetic "employee" begins routine activity:

- ✓ Normal logins during business hours
- ✓ Participation in Slack threads
- ✓ Pull requests submitted
- ✓ MFA configured

However:

- Authentication shows periodic consumer VPN patterns
- Minor geolocation inconsistencies
- Endpoint telemetry includes unfamiliar background processes

Individually, none are high-confidence indicators.
Signals exist. No escalation.

Day 31-45

PHASE 03

Privilege Expansion

Through legitimate collaboration

- ✓ Gains access to additional repositories
- ✓ Receives temporary elevated permissions for a project
- ✓ Obtains API credentials for a testing integration
- ✓ Accesses documentation of vendor-facing systems

No privilege escalation exploit. Access granted through workflow.
Trust compounds access.

Day 46-60

PHASE 04

Lateral Observation

The attacker begins mapping

- ✓ Internal SaaS integrations
- ✓ OAuth token flows
- ✓ Third-party vendor portals
- ✓ CI/CD pipelines

- Data exfiltration remains low-volume and staged.
- Telemetry shows: Legitimate use of authorised tools.
- No malware deployed. No alerts triggered.

Day 61-75

PHASE 05

Supply Chain Pivot

Using valid credentials and integration tokens:

- ✓ Accesses a shared vendor testing environment
- ✓ Interacts with API endpoints trusted by downstream partners
- ✓ Introduces subtle configuration modification

No ransomware.
No visible disruption.
But trust relationships are now exploited.

Day 76-90

PHASE 06

Monetisation / Exploit

Outcome varies depending on scenario:

Option A:
Sensitive data staged and sold quietly.

Option B:
Access resold

What makes insider-enabled fraud particularly consequential is not the access itself, it is what that access allows an attacker to do across two layers simultaneously.

External attackers typically operate on one surface at a time. They compromise an identity, or they exploit a payment vulnerability. Bridging the two requires additional steps, additional exposure, and additional risk of detection. An insider, synthetic or otherwise, already sits at the intersection. They hold identity credentials and operate within payment workflows as a matter of routine. The combination is not an escalation. It is the default.

The mechanic operates at a greater scale in financial institutions where insider positioning enables account takeover, fraudulent payment initiation, or the quiet redirection of transaction flows through accounts the insider controls or has been paid to overlook.

The structural implication is this: detection built around external threat models — perimeter breaches, anomalous login attempts, known malware signatures — will not catch this. **What is required is continuous, cross-domain visibility that treats identity behaviour and payment activity as a unified signal, monitored from onboarding through offboarding, with particular sensitivity to the moments where trust compounds access — exactly as it did on days 31, 45, and 61 of the taxonomy above.**

03

The Need For Continuous, End-to-End Monitoring Across the Customer Journey

PSR requires fraud monitoring that actually prevents and detects fraud, and fraud does not confine itself to the moment of authentication. This is not a point-in-time problem. A customer session in digital banking is a journey; it begins the moment a user initiates contact with the bank's digital surface and does not end at transaction authorisation. At every point in that journey, the fraud risk profile is evolving.

Under PSR, what happens between authentication and payment is where some of the most consequential fraud occurs. A fraud tool that fires a single API call at login misses everything that matters.

3.1 The Four Phases of PSR-ready Monitoring

PSR Article 83 does not ask PSPs to authenticate users once and consider the job done. It requires monitoring mechanisms to analyse environmental and behavioural characteristics throughout the payment journey — continuously, not episodically. End-to-end monitoring spans four distinct phases, each carrying its own fraud risk profile.

PHASE 1

Pre-Authentication

Before any credentials are entered, the session environment carries intelligence. Establish baseline risk before a single keystroke.

- Device fingerprint
- IP reputation
- VPN/proxy presence
- Prior mule device history
- Geolocation

PHASE 2

Authentication

SCA is applied and risk scored. Under PSR, this is the beginning of monitoring — not the end. TRA exemption must be supported by Phase 1 signals.

- SCA result
- Credential risk
- Behavioural biometrics
- TRA decision inputs

PHASE 3

In-Session Activity

Most underserved by legacy tools — most valuable to fraudsters. Post-authentication, the session environment can change in ways invisible to point-in-time tools.

- Remote access tool activation
- Screen-share detection
- Incoming call
- SIM-swap indicators
- Behavioural anomalies

PHASE 4

Payment Initiation & Post-Auth

Full accumulated signal picture from Phases 1–3 must inform the real-time risk decision. Payee PSP obligations also activate at this phase.

- Combined signal risk score
- Transaction context
- 10-second instant payment execution window (Instant Payments Regulation, EU 2024/886)

Each of these four phases generates signals. Each of those signals, under an API-based pricing model, represents a potential charge. The commercial incentive to limit signal collection is in direct tension with the regulatory requirement to maintain it continuously throughout all four phases.

Many fraudulent attempts stay undetected fraudulent attempts stay undetected for a long period of time while fraud and fraud rates are rising. Why? Because most systems might have recorded fraud, but did not stop it at earlier stages due to fragmental approach. Most of the cases are reported at later stages.

Stop fraud in realtime

Having to wait for an account to be confirmed Fraudulent before reporting has limited results

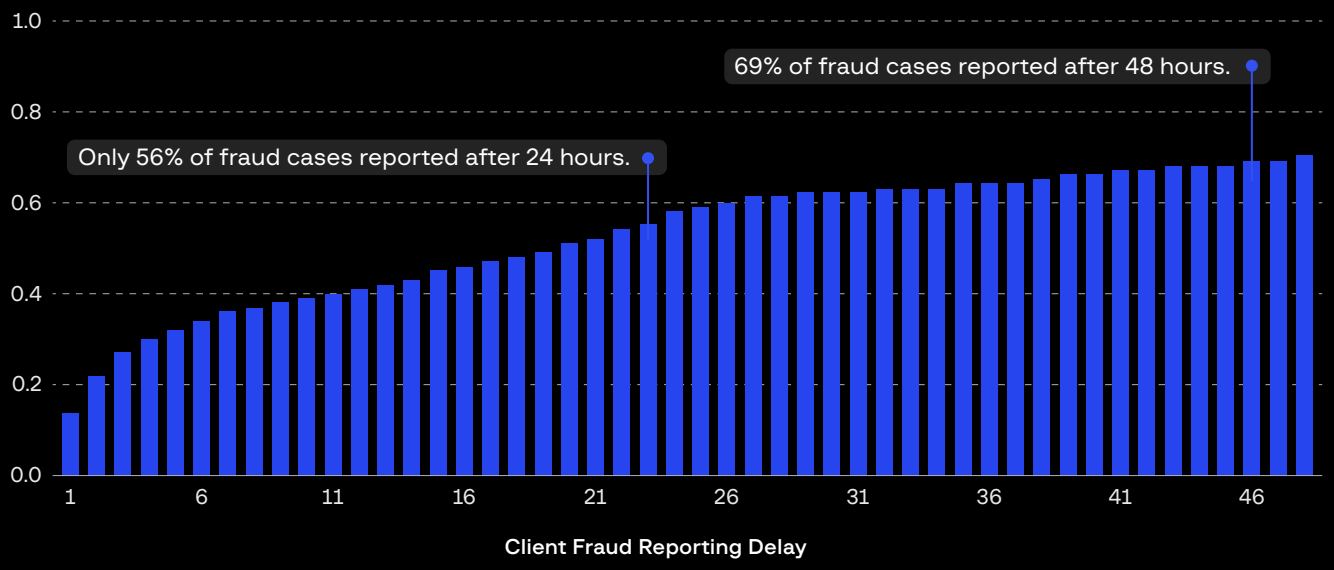

Image: Only 56% of fraud cases are reported within 24 hours; 69% only after 48 hours.

Why implement continuous monitoring?

Mule account creation example

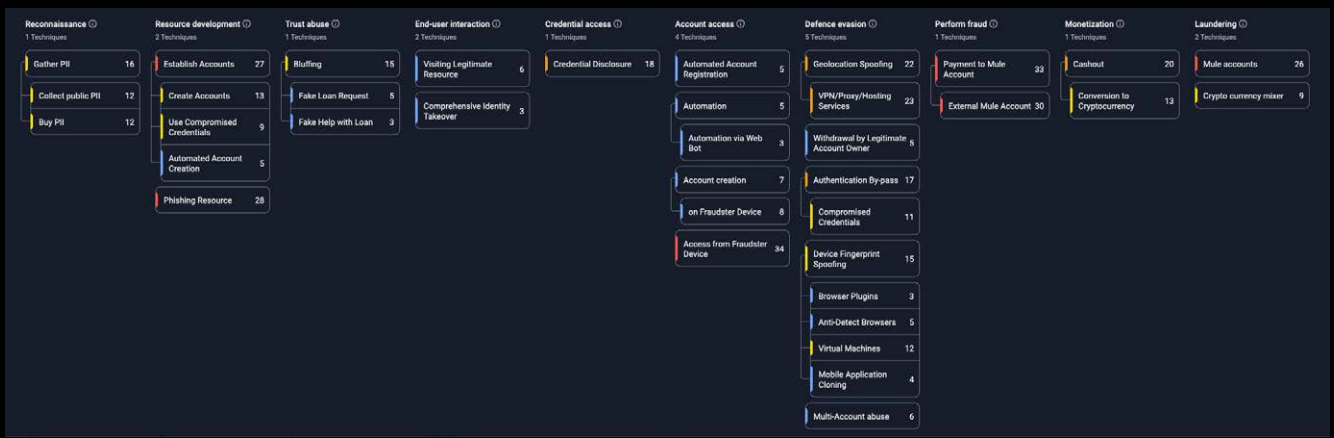

Image: Anatomy of a Fraud Operation: Mule Account Exploitation on French B2B Fintech Platforms

End-customer flow vs checks

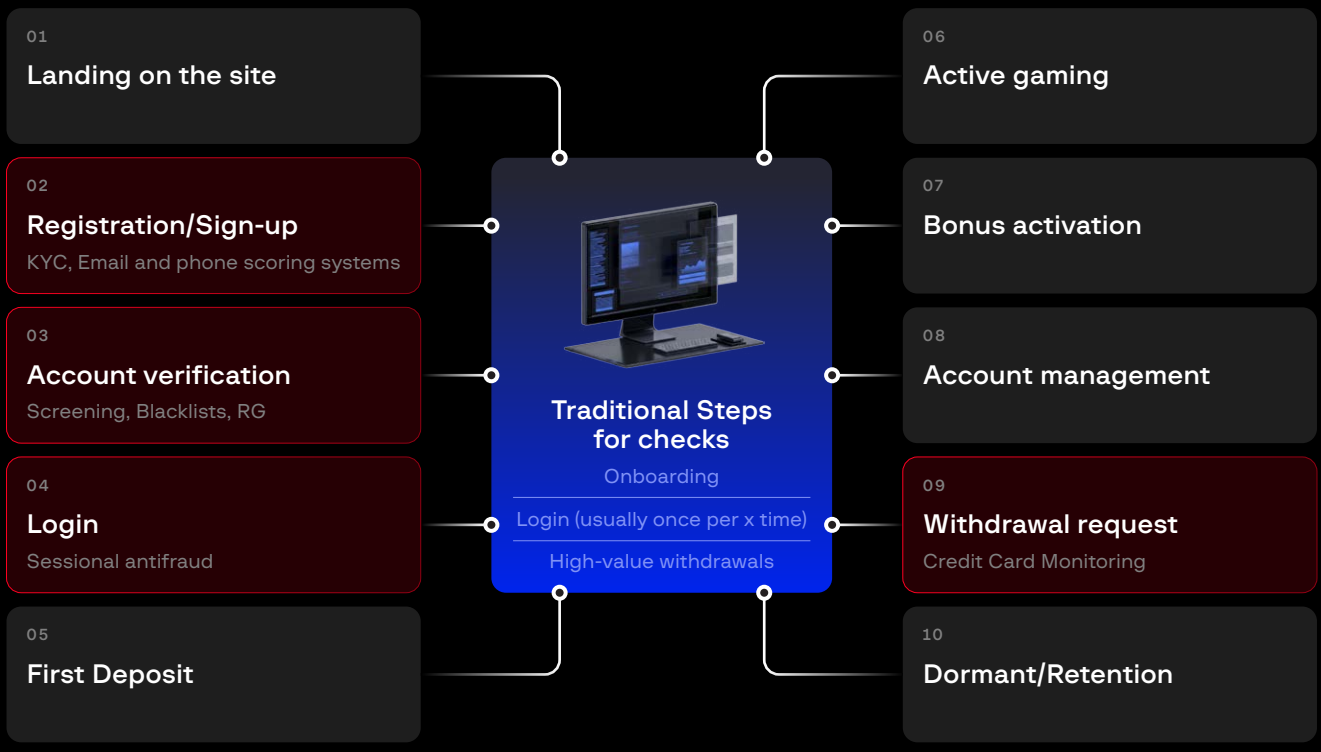

Image: Traditional way of performing checks is not sufficient to fulfill the PSR monitoring requirement.

04

Shifting to Proactive Fraud Prevention with Intelligence at the Core

Proactiveness isn't a setting. It's a posture. And most institutions are not there yet — not because they lack technology, but because they are measuring the wrong metrics: detection rates, chargeback ratios, fraud write-offs. These are all outcomes of events that have already happened. They tell you how well you responded. They tell you nothing about whether you could have stopped the attack from reaching the transaction at all. The metrics that matter in building a proactive posture are not outcome metrics; they are readiness and anticipation metrics.

Here is what proactive actually means at Group-IB: visibility into adversary infrastructure being staged before an attack launches. Fraud kits being assembled, mule networks being recruited, synthetic identities being tested: the preparation phase that precedes every large-scale fraud campaign. **In 2025, Group-IB's operations detected over 1,800 active phishing kits, flagged 233% more darknet sellers of attack tooling year-on-year, and identified infrastructure staging for campaigns targeting European financial institutions an average of five hours before activation. Those are readiness metrics. Most institutions have none of them on their dashboard.**

That visibility does not come from transaction data. It comes from intelligence that originates on the other side of the operation. Group-IB has been inside criminal ecosystems, not as observers, but as active investigators working alongside Interpol, Europol, Afripol and national cybercrime units across the globe.

4.1 Building Readiness for Safe and Compliant Intelligence Sharing

Regulators across Europe, the UK, Singapore, Australia, and the US have reached the same conclusion: institutions must share fraud intelligence collectively or keep losing individually. The requirements exist. The political will is there. The problem has never been whether to share, but how to share suspicious activity in real time without violating the privacy laws that govern every institution in the network.

This is the Intelligence-Sharing Paradox, and it has a precise cost: RUSI research shows 28% of fraudulent funds leave a mule account within 15 minutes. The traditional confirmed-fraud model shares that intelligence 10 to 21 days later. We are fighting 15-minute crimes with 24-hour data.

The technical barrier is equally specific. Standard SHA-256 hashing — the industry default — is vulnerable to dictionary attacks. As payment data standardises under ISO 20022, those dictionaries become universal. Hashing does not meet Data Protection by Design. It never did.

Group-IB's patented Distributed Tokenisation breaks a sensitive identifier into parts, distributes them across participating entities for tokenisation, and ensures raw personal data never leaves the originating institution. The result is a consistent, reversible-by-no-one token that works for network-wide pattern recognition. Bureau Veritas confirmed GDPR compliance independently in September 2025.

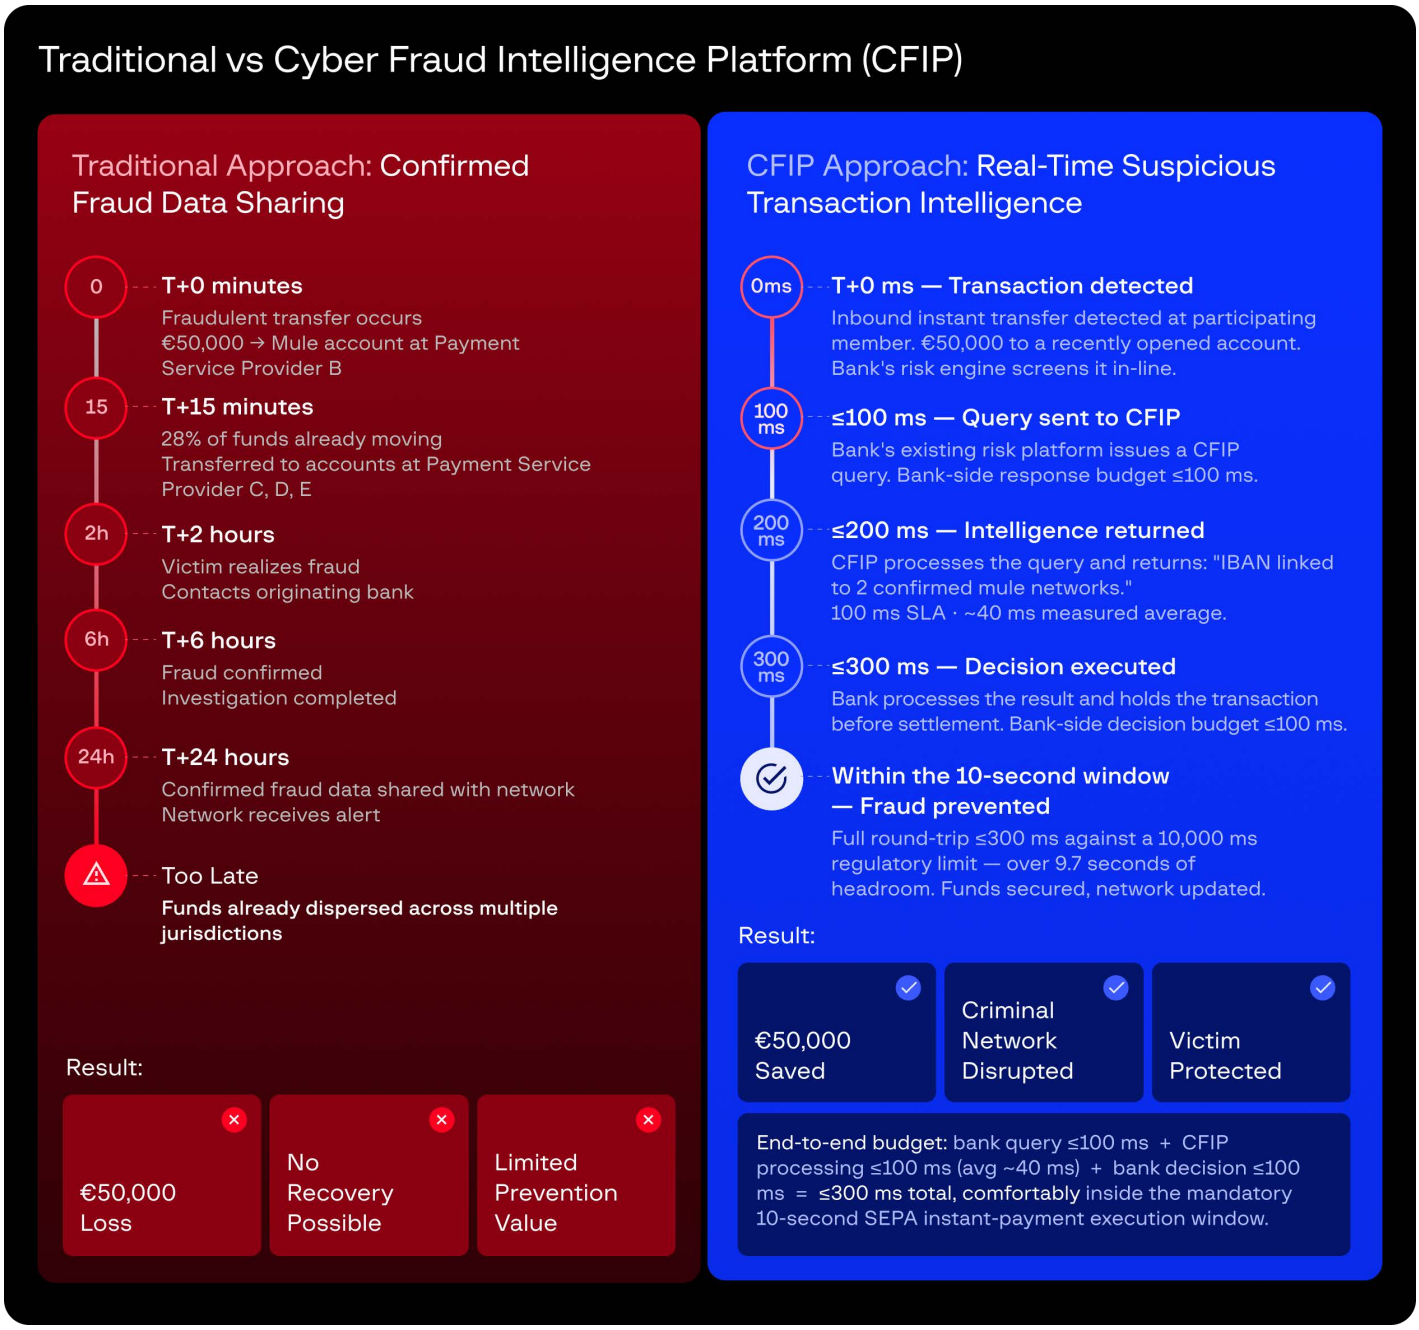

Cyber-Fraud Intelligence Platform (CFIP) bridges one of the biggest gaps in financial-crime prevention — the ability to collaborate securely. By combining real-time risk-signal sharing with independently verified privacy safeguards, CFIP allows the industry to work together to stop fraud before it happens.

4.2 Frameworks and Ecosystems that Enable Intelligence Sharing

The proof that this works is not theoretical. A Central Bank in Central Asia deployed CFIP as a national intelligence hub, connecting participating financial institutions under a single real-time layer. With just 5% of member banks performing checks, the network identified 300–400 mule accounts daily, processed 180,000 checks per day, and prevented an estimated $10–15 million in fraud annually. At full adoption, projections reach $100–300 million.

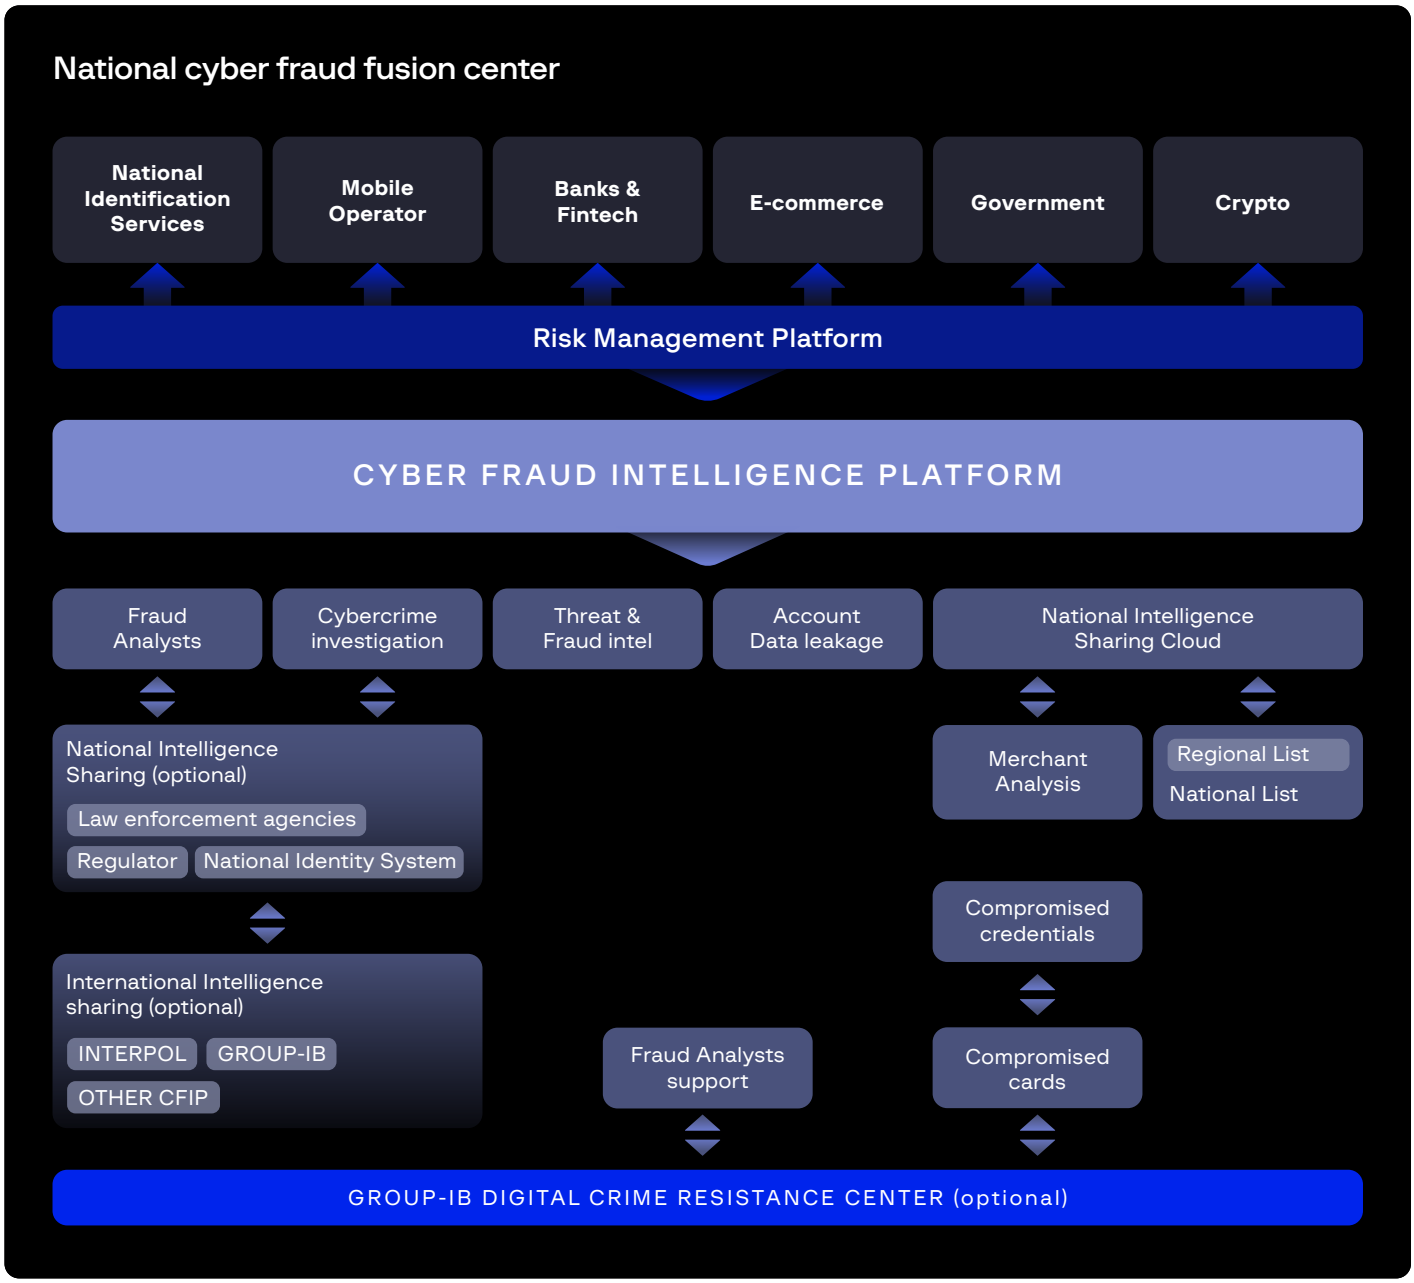

The architecture scales from single institutions to national hubs to cross-border networks — without replacing any existing fraud infrastructure. CFIP operates as an intelligence layer on top of what institutions already have. When an internal risk engine flags a transaction, CFIP queries the network on that specific signal. No data lakes. No broad mining. Purpose limitation and data minimisation are satisfied by design, not policy.

The network effect is the point. Every flagged mule account during the warm-up phase forces criminal networks to restart an expensive

account-conditioning process. Every participating institution makes every other one stronger. Fraudsters already operate this way as coordinated, cross-institutional networks. The only effective counter is a defence built the same way.

Collective defence requires shared language. Group-IB Fraud Matrix provides it — a MITRE ATT&CK-inspired framework built into the Unified Risk Platform that maps the full cyber-fraud kill chain: from reconnaissance and account conditioning through defence evasion and monetisation.

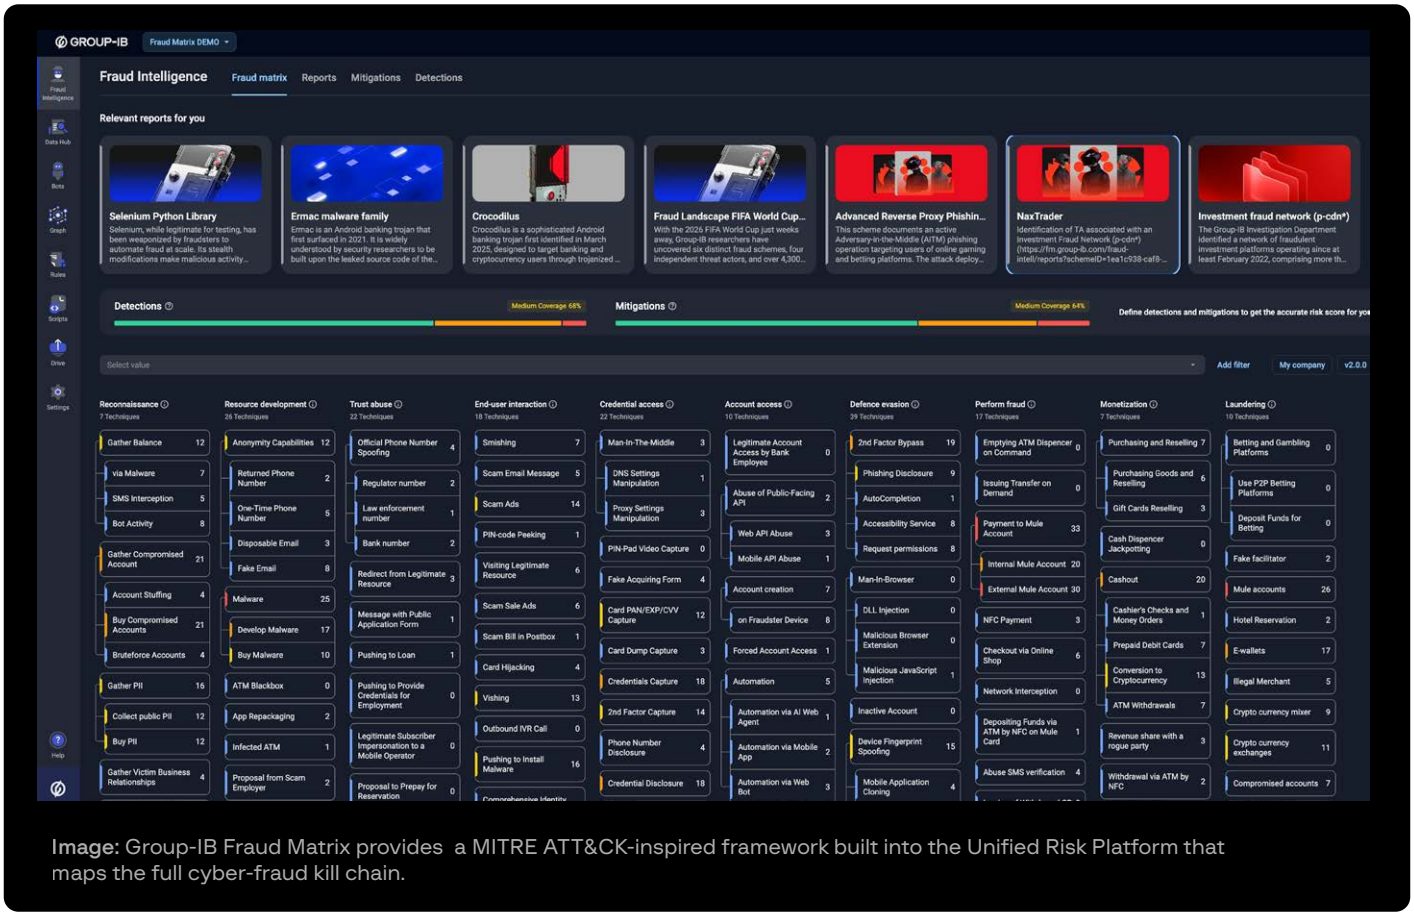

Image: Group-IB Fraud Matrix provides a MITRE ATT&CK-inspired framework built into the Unified Risk Platform that maps the full cyber-fraud kill chain.

What makes it operationally relevant rather than taxonomic is the data underneath it. Every technique is mapped to observed frequency from Group-IB's intelligence operations — not modelled risk, but actual attack patterns.

When deepfake KYC bypass attempts grew 80% in the final four months of their tracked period, that signal propagated into CFIP's network-wide detection thresholds before most institutions had encountered it individually. **That is the connection between the Fraud Matrix and Cyber Fraud Intelligence Platform: the former tells you what adversaries are doing; the latter ensures the network acts on it before you become the next data point.**

Upcoming releases will extend the framework to full TTPs including tooling and procedures, implementation guidance for countermeasures, and localised detection thresholds calibrated to regional regulatory requirements.

05

Lower Operational Costs Through Automation and Higher Alert Precision

The regulatory shift across Europe fundamentally changes the economics of financial crime defense. Under the Payment Services Regulation (PSR) Article 83 requirement, selective or transaction-triggered fraud monitoring will fall short of what PSR requires. Payment Service Providers (PSPs) are legally required to implement continuous, session-wide instrumentation—covering device intelligence, behavioral biometrics, remote access detection, and malware presence—across every single customer journey.

However, this requirement creates a dangerous economic trap under traditional event-based or per-API pricing models. The more thoroughly an institution monitors its sessions to comply with the law, the more API calls it fires, causing vendor bills to grow in direct proportion to compliance thoroughness. This introduces a perverse financial incentive to ration data collection and compromise security to control operational costs.

An intelligence-led, flat-rate pricing model completely resolves this tension by charging the PSP a single annual fee calibrated to the size of its active customer base, rather than the volume of sessions, API calls, or fraud events generated. This commercial structure wholly decouples the depth of risk monitoring from the ultimate invoice, making it the only commercial framework structurally compatible with PSR-ready fraud prevention at scale.

The Strategic Advantages of the Flat-Rate Model

- **Total Cost Predictability:** The PSP knows its annual fraud prevention spend on day one. The invoice remains unchanged regardless of mid-year instant payment volume surges, intensive fraud campaigns, or regulatory upgrades demanding enhanced session monitoring. For finance and risk teams building multi-year compliance budgets under evolving European Banking Authority (EBA) standards, this predictability is an operational prerequisite.
- **No Incentive to Ration Context:** Because annual costs are fixed, the PSP has zero financial reason to limit signal collection. Fraud teams can instrument every phase of the customer session as richly as the threat landscape demands, eliminating the structural blind spots forced by per-event billing.

- **Continuous In-Session Monitoring:** The in-session phase—the critical window between authentication and payment initiation—is often the most commercially prohibitive to monitor under event-based pricing due to its high signal density. A flat-rate model makes comprehensive, continuous remote access tool (RAT) detection, call-intercept monitoring, and behavioral anomaly analysis economically viable throughout the entire journey.
- **Audit Trail Completeness for Liability Defense:** The PSR liability framework shifts the burden of proof to the bank, requiring institutions to demonstrate that robust fraud prevention mechanisms were active at the exact time of a disputed transaction. Flat-rate pricing removes the cost disincentive to logging complete sessions, generating the continuous audit trail necessary to defend against reimbursement claims and regulatory reviews.
- **False Positive Reduction and TRA Exemption Retention:** Access to richer, unrationed contextual signals improves machine-learning precision, allowing systems to accurately distinguish genuine users from fraudsters. This directly reduces false positive rates, minimizes customer friction, and ensures the bank maintains the low fraud thresholds required to retain its Transaction Risk Analysis (TRA) exemption for frictionless checkouts.

Quantitative and Qualitative ROI Impact

Transitioning to an architecture anchored by flat-rate predictability and rich, unrationed telemetry yields measurable balance-sheet performance across three distinct vectors.

ROI Vector	Operational & Financial Impact
Loss Reduction	Financial institutions implementing predictive analytics with unrationed device behavioral signals have demonstrated up to a **60% reduction in total fraud losses.**
Operational Efficiency	Higher alert precision cuts **false positives by 50%**, significantly reducing manual analyst investigation workloads and accelerating the mean time to resolve (MTTR) active alerts
Customer Lifetime Value	Precision authentication applies friction only when high risk warrants it. Minimizing unnecessary step-up challenges improves transaction conversion rates, prevents customer churn, and strengthens brand trust.

60%

Fraud loss reduction achievable

With predictive analytics + rich signals

50%

Reduction in false positives

Improving customer experience

0

Incremental cost for deeper monitoring

Under flat-rate pricing model

5.1 Qualitative ROI: Superior Safety as a Competitive Differentiator

In a market where PSR requirements are creating a minimum safety floor, competitive differentiation will be determined by who exceeds that floor with confidence and efficiency — not just who meets it.

Institutions that build genuine intelligence-led fraud prevention capability will hold advantages across three vectors: customer trust (measurably lower fraud rates become a marketable trust signal), regulatory standing (proactive, comprehensive monitoring generates the evidence base that regulatory examinations require), and partnership value (institutions operating at the frontier of collective defence attract stronger ecosystem partnerships and preferential regulatory engagement).

Ultimately, the shift from reactive to proactive fraud prevention is not a compliance exercise. It is a structural repositioning from fraud management to fraud prevention — with every element of the intelligence stack aligned to stopping fraud before it reaches the transaction, rather than responding after funds have moved.

Conclusion: Intelligence Before the Transaction

From Isolated Events to Disrupted Campaigns

European financial institutions have spent the last decade building faster reactions. **The next decade will be won by those who build earlier visibility: before the fraud kit is assembled, before the mule account is recruited, before the synthetic identity reaches your onboarding queue. And increasingly, before an autonomous agent completes the attack cycle without a human ever making a decision.**

That shift is not theoretical. Group-IB's intelligence infrastructure monitors the precise signals that precede every large-scale fraud campaign — underground market activity, infrastructure staging, credential exposure, account conditioning patterns, SIM pre-event signals — detected an average of five hours before activation.

The threat this infrastructure monitors has also evolved. The adversary Group-IB tracks in 2026 is not always a criminal network optimising for profit. In several documented cases it is a state-aligned operation using financial fraud as cover, funding mechanism, or strategic disruption tool. The intelligence required to detect both is the same. The architecture that stops a coordinated criminal campaign is the same architecture that stops a geopolitically motivated one.

These are not post-transaction flags. They are pre-crime warnings.

In 2025 alone, Group-IB's operations conducted through 11+ Digital Crime Resistance Centres across 60 countries and in direct partnership with Interpol, Europol, Afripol and national cybercrime agencies resulted in 1,809 arrests. That is not events detected. That is coordinated campaigns disrupted at the source.

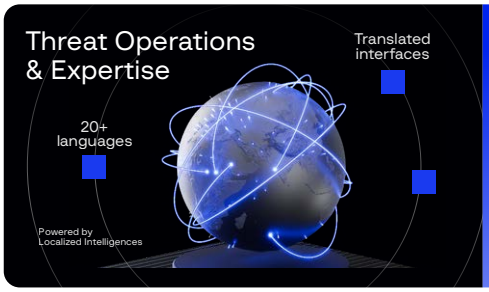

LOCALIZED PRESENCE AND INTELLIGENCE INFRASTRUCTURE

Singapore, Vietnam, Thailand, Malaysia, Egypt, Netherlands, Italy, Chile, Uzbekistan, UAE, KSA

Power by Unique Digital Crime Resistance Centres

For fraud and security leaders in European financial institutions, the question is no longer whether proactive intelligence-led prevention is possible. It is whether your current architecture, your tooling, your team structures, your commercial agreements with vendors, is built to support it. Under PSR, the cost of that answer being "no" now lands directly on your balance sheet.

Build Your Edge Against Fraud With Group-IB

Request a Threat Exposure Assessment

Understand exactly how your institution appears from the adversary side — which credentials are circulating on underground markets, whether your infrastructure or brand is being staged for an upcoming campaign, and where your customers are currently most exposed. Delivered by Group-IB analysts with direct access to criminal ecosystem intelligence.

BOOK A BRIEFING SESSION

Assess Your PSR Readiness

PSR's Article 83 continuous monitoring requirement is not a future concern, it is an immediate compliance risk for any PSP relying on selective, API-based fraud instrumentation. Group-IB's PSR Readiness Assessment evaluates your current monitoring architecture against the regulation's four-phase session requirements and delivers a prioritised remediation roadmap within two weeks.

START YOUR ASSESSMENT

See the Cyber Fraud Intelligence Platform (CFIP) in Action

Group-IB's Cyber Fraud Intelligence Platform is the only GDPR-validated, privacy-preserving solution that enables real-time cross-institutional fraud signal sharing before fraud is confirmed. In a 45-minute working session, we will show you how it maps to your PSR obligations, what signals it would have flagged in your last three major fraud incidents, and what onboarding looks like for a European PSP.

SCHEDULE A DEMONSTRATION

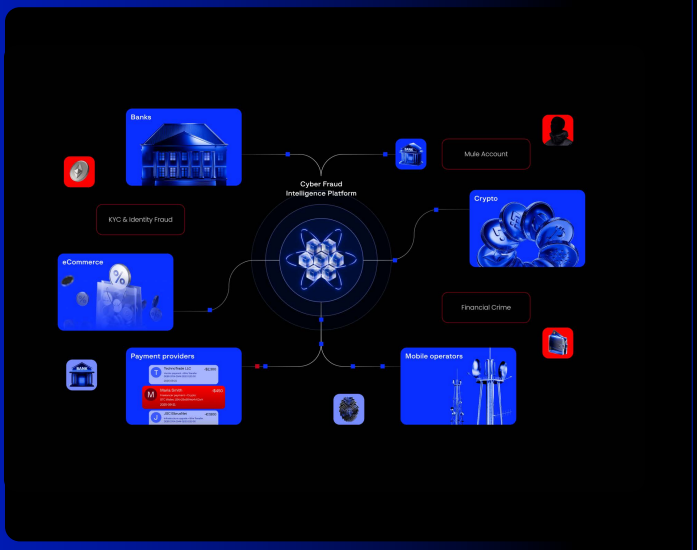

Endnotes

1. The Voice of Fraud: Deepfake Vishing and the New Age of Social Engineering
2. High-Tech Crime Trends Report 2026
3. Fraud Detection in Banking: 2025 Future Trends & Predictions
4. Fraud Detection in Banking: 2025 Future Trends & Predictions
5. Weaponized AI: Inside The Criminal Ecosystem Fueling The Fifth Wave of Cybercrime
6. The Voice of Fraud: Deepfake Vishing and the New Age of Social Engineering
7. Weaponized AI: Inside The Criminal Ecosystem Fueling The Fifth Wave of Cybercrime
8. Fraud Detection in Financial Services Using Advanced Machine Learning
9. Fraud Detection in Banking: 2025 Future Trends & Predictions
10. High-Tech Crime Trends Report 2026
11. Cyber Saga: In the Footsteps of the DPRK IT Workers
12. Fraud Detection in Banking: 2025 Future Trends & Predictions
13. AML Fraud Detection: The Hidden Threats Banks Miss in 2025
14. AML Fraud Detection: The Hidden Threats Banks Miss in 2025
15. Authorized Push Payment (APP) Fraud: A Global Threat

GROUP-IB

Group-IB is a creator of cybersecurity technologies to investigate, prevent and fight digital crime.

1,600+
Successful investigations of high-tech crime cases

500+
Employees

60
Countries

$1 bln+
Saved by our client companies through our technologies

#1*
Incident Response Retainer vendor
*According to Cybersecurity Excellence Awards

11
Unique Digital Crime Resistance Centers

Global partnerships

- INTERPOL
- EUROPOL
- AFRIPOL

Recognized by top industry experts

FORRESTER® · datos INSIGHTS · kuppingercole ANALYSTS · Gartner · IDC · FROST & SULLIVAN

Fight against cybercrime

GROUP-IB.COM
INFO@GROUP-IB.COM

APAC
+65 3159 4398

EU & NA
+31 20 226 90 90

MEA
+971 4568 1785