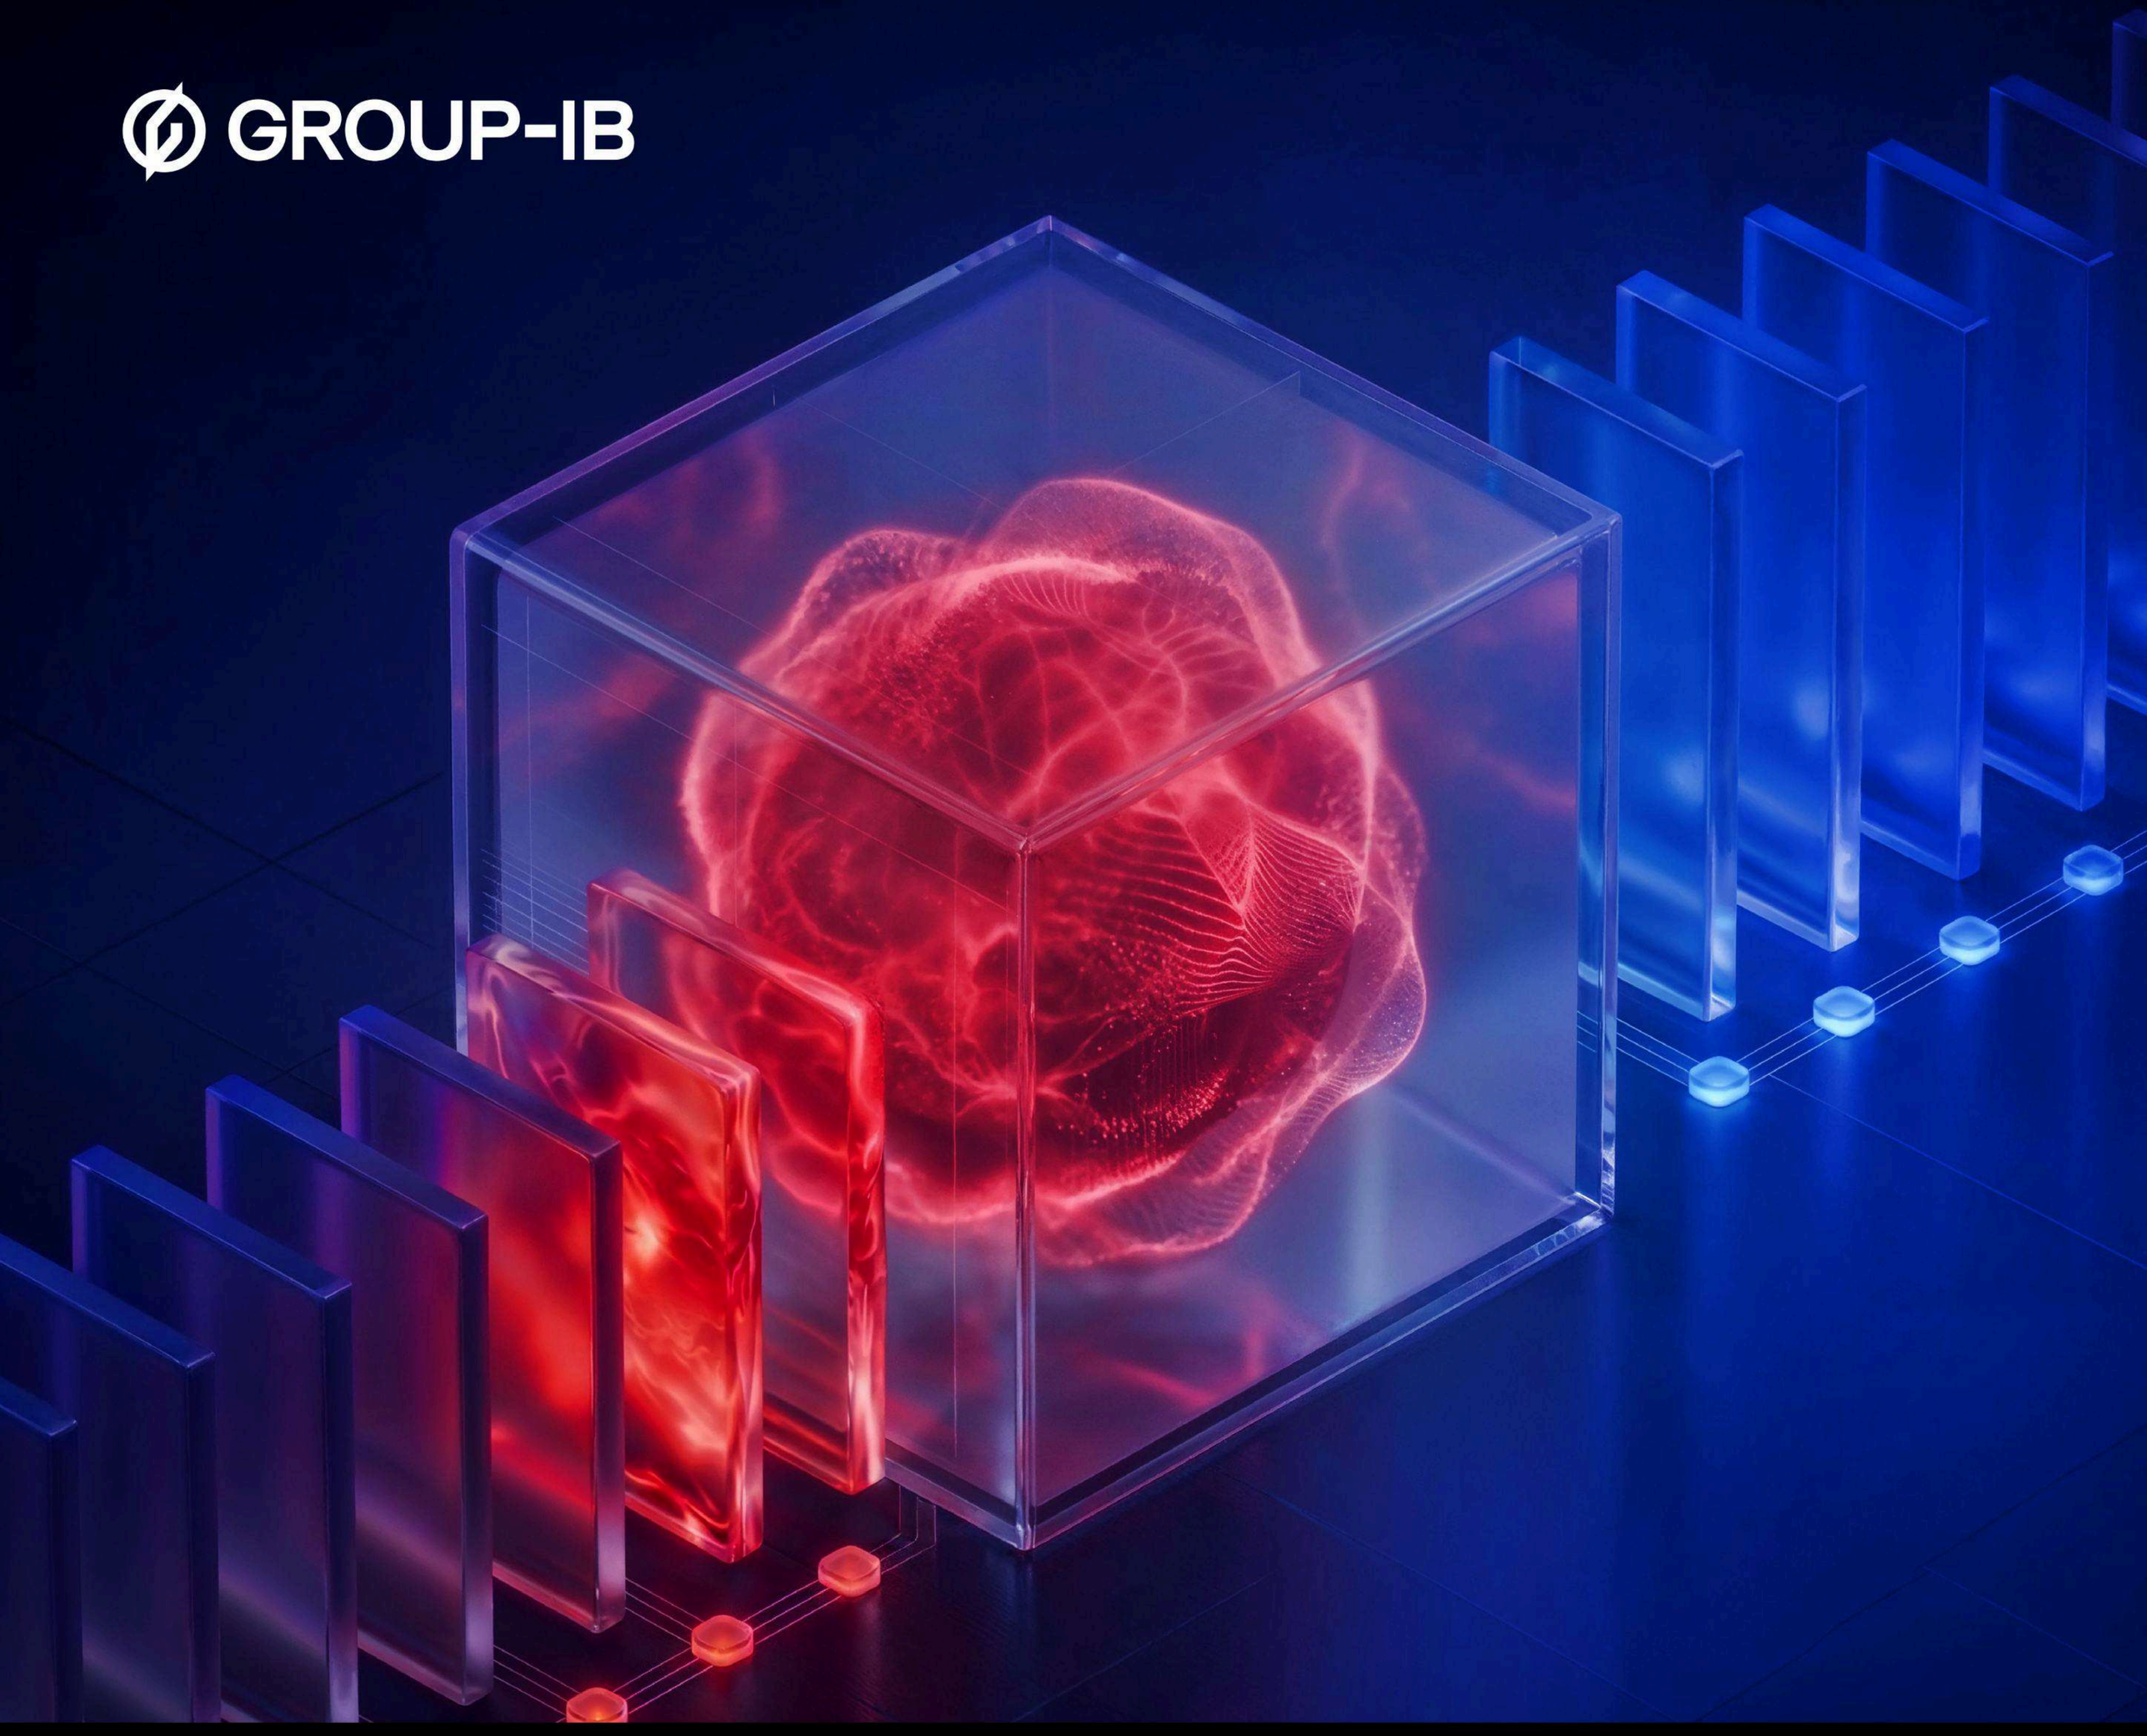


White Paper

Before Fraud Transacts

Enabling Proactive Prevention in the Age of Industrialized AI

GROUP-IB.COM

Table of contents

Executive Summary 03

01 The Threat 06

02 The Gap 10

03 The Operating Model 17

04 Proof 23

05 What This Changes 24

The Decision 29

Technical Annex · For Enterprise Architects 31

Executive Summary

The case in one page. The blueprint in the next twenty.

Over the past 18 months, a recurring theme has emerged in conversations with fraud and risk leaders: "We have invested in fraud detection, threat intelligence, and KYC platforms. Each delivers value in its own domain, yet we still struggle to identify coordinated campaigns early enough to prevent them.

This paper is not a product pitch. It is a diagnostic framework for understanding why that gap persists and a practical operating model for closing it. While our approach reflects our experience building solutions in this space, the framework itself is intended to contribute to the broader industry conversation.

Fraud Has Become a Campaign, Prevention Must Too

The adversaries targeting financial institutions today do not exploit individual transactions. They run coordinated campaigns with supply chains, operational timelines, infrastructure networks, and role-specialised teams that treat each transaction as the final stage of a process begun weeks or months earlier.

The transaction is not where fraud happens. It is where fraud ends. This distinction between defending events and disrupting campaigns is the single most important strategic shift in financial crime prevention today.

The market is consolidating around platforms that can prove an outcome: closed-loop enforcement, fraud, and threat intelligence fused at the source, attribution with real depth behind it. Passive feed consumption, i.e., however polished the dashboard, it is quietly becoming indefensible as a standalone strategy. This paper argues that the path forward is prediction-first defence with cyber-fraud fusion: the convergence of cybersecurity intelligence, fraud prevention, into one operating model with the main objective of retrieving intelligence before the transaction.

This is not Group-IB's prediction alone. In Gartner's first-ever Magic Quadrant™ for Cyberthreat Intelligence Technologies (2026), Group-IB was named a Leader — one of five, out of eighteen vendors evaluated. Gartner cited the Cyber Fraud Fusion model directly, both as a current strength and in a roadmap it described as a shift "from detection to prediction-first defence." The category is reorganising around exactly this, and the institutions that survive the decade will be the ones that move from observability to consequence before the move is forced on them.

The proposed operating model is not theoretical. It is aligned with the MITRE Fight Fraud Framework™, the industry's first standardised fraud kill-chain taxonomy. Its cross-institutional layer, CFIP, runs on patented distributed tokenisation **independently validated as GDPR-compliant (Article 26) by Bureau Veritas**. The architecture has been deployed across financial institutions in multiple jurisdictions and has contributed to over US$100M in prevented losses in 2025.

$579.4 billion

lost to bank fraud and scams in 2025

Criminal networks are now operating with the scale and coordination of multinational corporations.

Nasdaq Verafin 2026 Global Financial Crime Report

Executive Takeaways — tl;dr

Stop defending events	The transaction is the final stage of a campaign already in motion. Your controls are positioned at the wrong point in the adversary timeline.
Break the silos	Telemetry collection and network queries run asynchronously alongside transaction processing, so the fraud layer does not add latency to the payment path itself.
Converge before the mandate forces you to	PSR, PSD3, Singapore SRF, Australia SPF, regulators in every major market are arriving at the same conclusion. If fraud passed through your institution, you are accountable. The evidence base to prove you tried to stop it must be built before the incident, not reconstructed after.

The Model In Three Layers

Cyber Fraud Fusion (CFF) is the operating model. It runs three capabilities as one closed loop

The same cybercriminal operation that steals credentials on Monday sells them on Wednesday, automates account takeover on Friday, and routes the proceeds through a mule network the following week. The attack is integrated. The supply chain is fused. The adversary has been running one process for years.

Most institutional defences were built to track disintegrated activities, so organisations built two separate functions: a cybersecurity team to watch the perimeter and a fraud team to watch the transaction stream.

Each function is excellent at what it does. Neither can see what the other sees. And the adversary, who operates across both, exploits that gap deliberately.

Cyber Fraud Fusion is where you stop running fraud, cyber, AML, as separate programmes, and to run them as one intelligence layer instead.

THE CYBER FRAUD FUSION MODEL

One process. Three layers. Across the adversary's timeline.

Each layer acts at a different point — before, during, and after the attack.

When it acts	Capability and what it does	
Pre-attack Weeks before the transaction	 Threat Intelligence + Digital Risk Protection	Detects attacker infrastructure, harvested credentials, and phishing kits before they reach customers.
In-session As the customer transacts	 Fraud Protection (FP)	Intercepts fraud at the session, not the transaction - through device, behavioural, and scam-call signals.
Post-attack As money moves between banks	 Cyber Fraud Intelligence Platform (CFIP)	Shares suspicions across institutions in real time, without exposing raw customer data.

What the model changes — in three numbers

The same $1.8M campaign, run through a siloed architecture versus a fused one.

When it acts	Today		With Cyber Fraud Fusion	
Malware to fraud-team awareness	23 days	architectural default	< 15 min	automated correlation
Suspicion to peer-bank action	10-21 days	manual SAR filing	< 5 min	cryptographic query
Financial outcome same campaign	$1.8M lost		> 90% prevented	

The numbers shown are based on a 46-banks pilot study

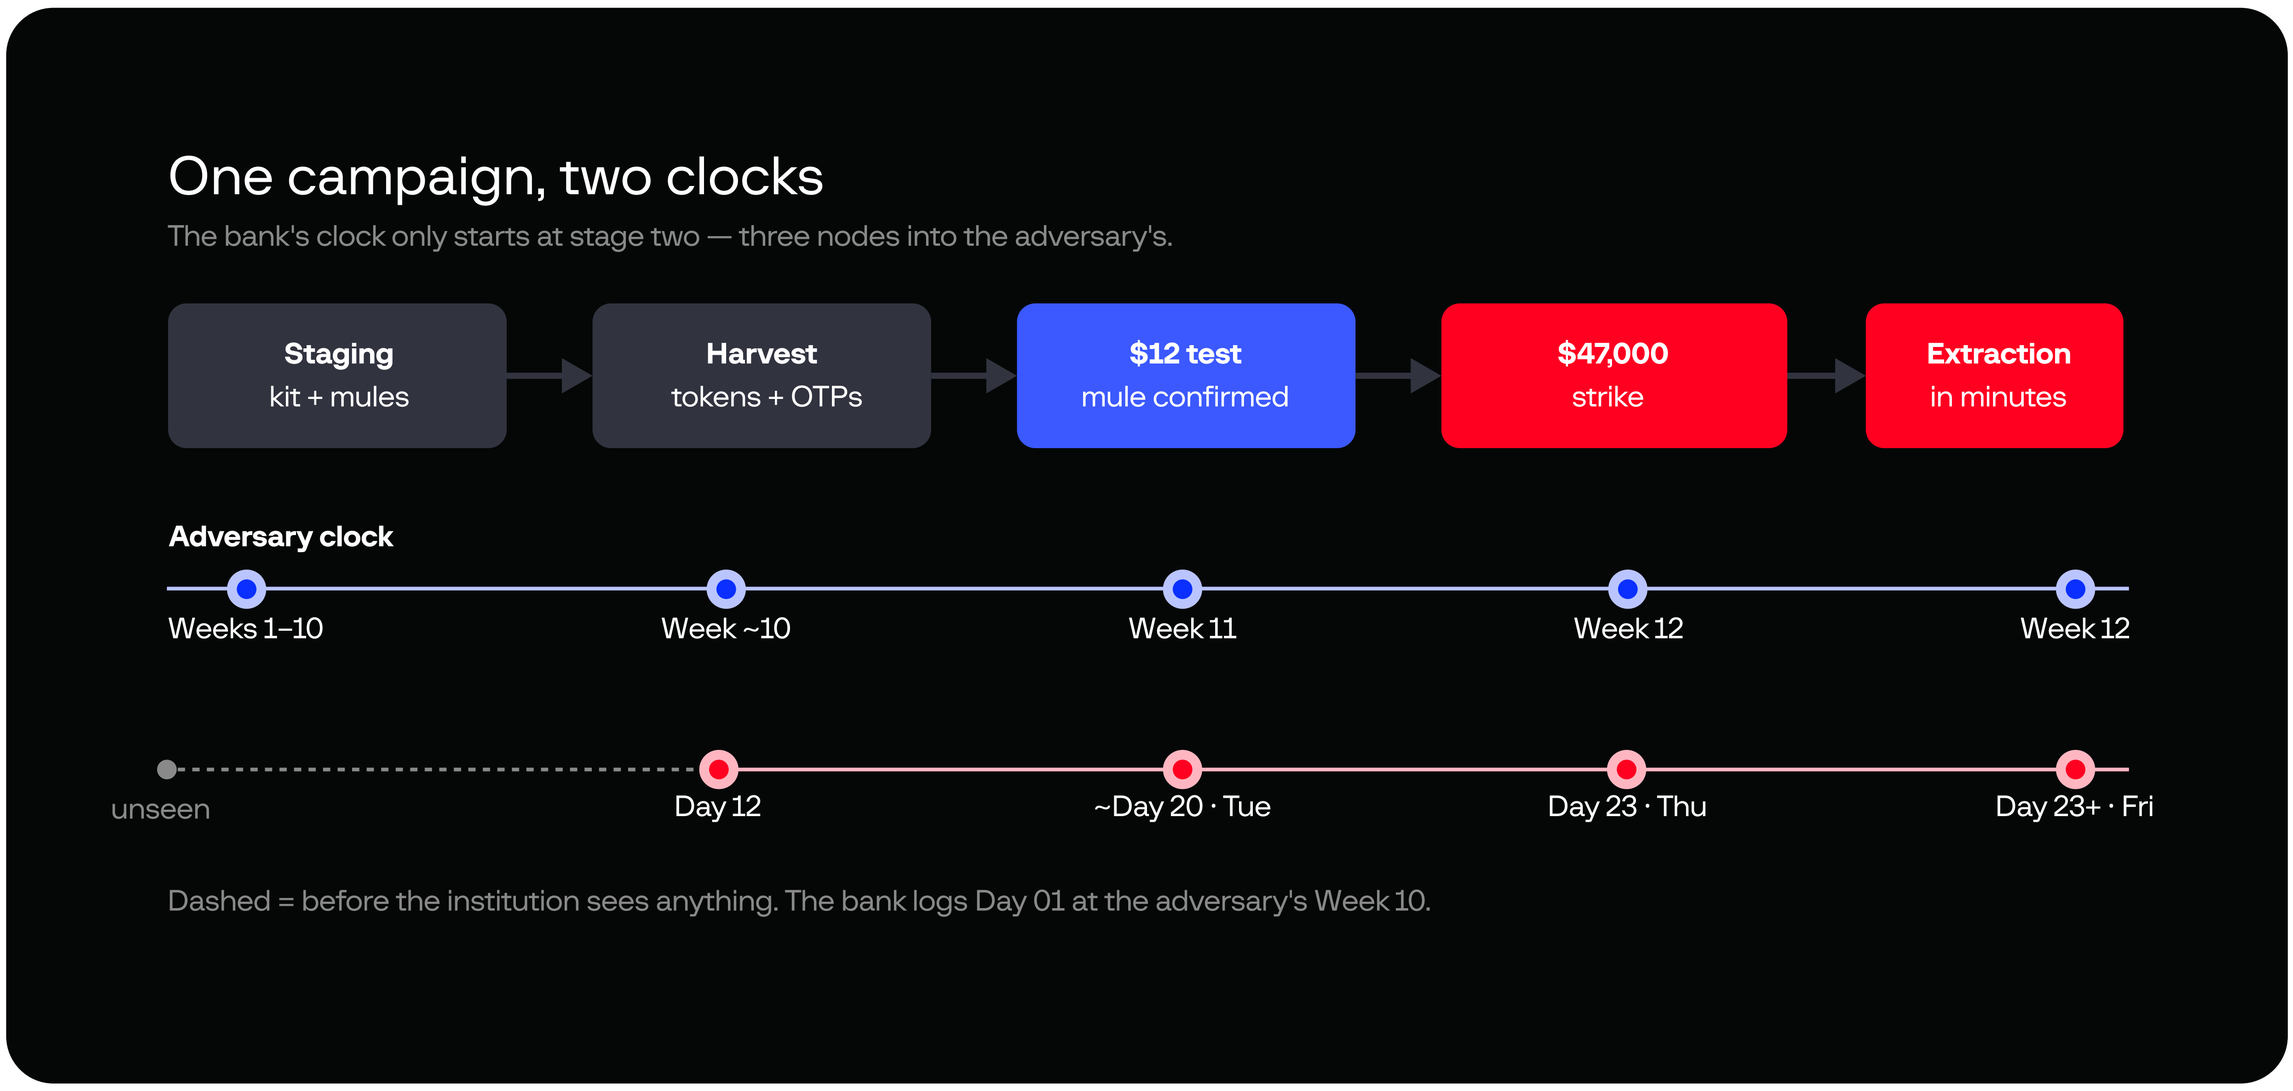
01

The Threat

The Campaign: 90 Days Before the Transaction

At 02:17 on a Tuesday morning, a mule account receives a low-value inbound transfer: $12. Nothing about it is unusual. The account is three months old, passes every behavioural baseline, and belongs to a verified customer. The receiving institution logs it and moves on.

By Thursday, that account will receive $47,000. By Friday, the money is gone. The $12 was not fraud. It was a test, a calibration. The campaign had been running for weeks. The institution had no idea. This is industrialisation in practice, not smash-and-grab, but strategy.



What follows is what those twelve weeks looked like

The supply chain that replaced the independent criminal

Modern financial fraud operates as a vertically integrated criminal supply chain: identity manufacturers, infrastructure operators, access brokers, social-engineering platforms, mule-network operators, and laundering specialists. Each role is a traded service. Each is replaceable. The supply chain is resilient by design, and built around the same logic as any legitimate business: specialise, automate, and optimise for margin.

The economics have inverted. High-volume campaigns generate high detection rates and diminishing returns. Precision campaigns invest in pre-campaign intelligence, calibrate to a specific institution's known controls, and rotate infrastructure before detection cycles complete.

Group-IB's investigation of the GXC Team, dismantled in Operation Big Bang, found more than 250 phishing sites and nine Android malware strains running simultaneously across Europe and the Americas as a subscription Crime-as-a-Service product. The same kit. Rotated. Refined. Reused against your peers before it reached you. The campaign does not end when one institution detects it. It graduates.

> The supply chain is now automated, language-agnostic, and self-optimising. But there is still one checkpoint the adversary cannot bypass with AI alone: the identity itself. If the identity is verified, the attack is trusted. So the adversary went further upstream — and made the identity.

The same identity pipeline, different expressions

The identity weakness that lets a fabricated employee through HR is the same weakness that lets a fabricated customer through onboarding. Both exploit one assumption: that a verified identity, once cleared, stays trustworthy. It does not.



Consider the engineer who had been with the firm for three months. Full system credentials, access to internal payment infrastructure, positive performance reviews, a LinkedIn profile with eight years of history.

Except, he had never existed.

Within weeks he requested elevated access to payment-processing systems for a project that turned out not to exist. Months later, investigators identified the employee as an AI-generated persona deployed by a state-backed actor. The resume was assembled from stolen identities, the video interviews conducted with real-time deepfake technology.

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

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.

8,065

Deepfake KYC bypass attempts documented by Group-IB between January–August 2025, with 80% concentrated in the final four months. Not a stable threat curve. An acceleration signal.

300+ organisations were infiltrated by AI-generated synthetic employees in 2025. The underground price for a full impersonation toolkit: $5. The same pipeline that harvested real data, localised device profiles, socially engineered liveness checks, immediate handover to fraud networks, also produces the mule accounts that pass KYC and behave normally for months before their first fraudulent transfer. The architecture that scores a deposit in isolation will never see the conditioning that made it dangerous.

The defence is visibility before the identity is trusted, not a smarter decision at the payment.

When the identity is compromised, every downstream system that trusts it is compromised too. But there is one final expression of this problem — the case where identity verification works correctly, the customer is genuine, and fraud still succeeds. It is authorised fraud, and it is where the money and the liability sit.

The hardest fraud to stop is the one your customer approves

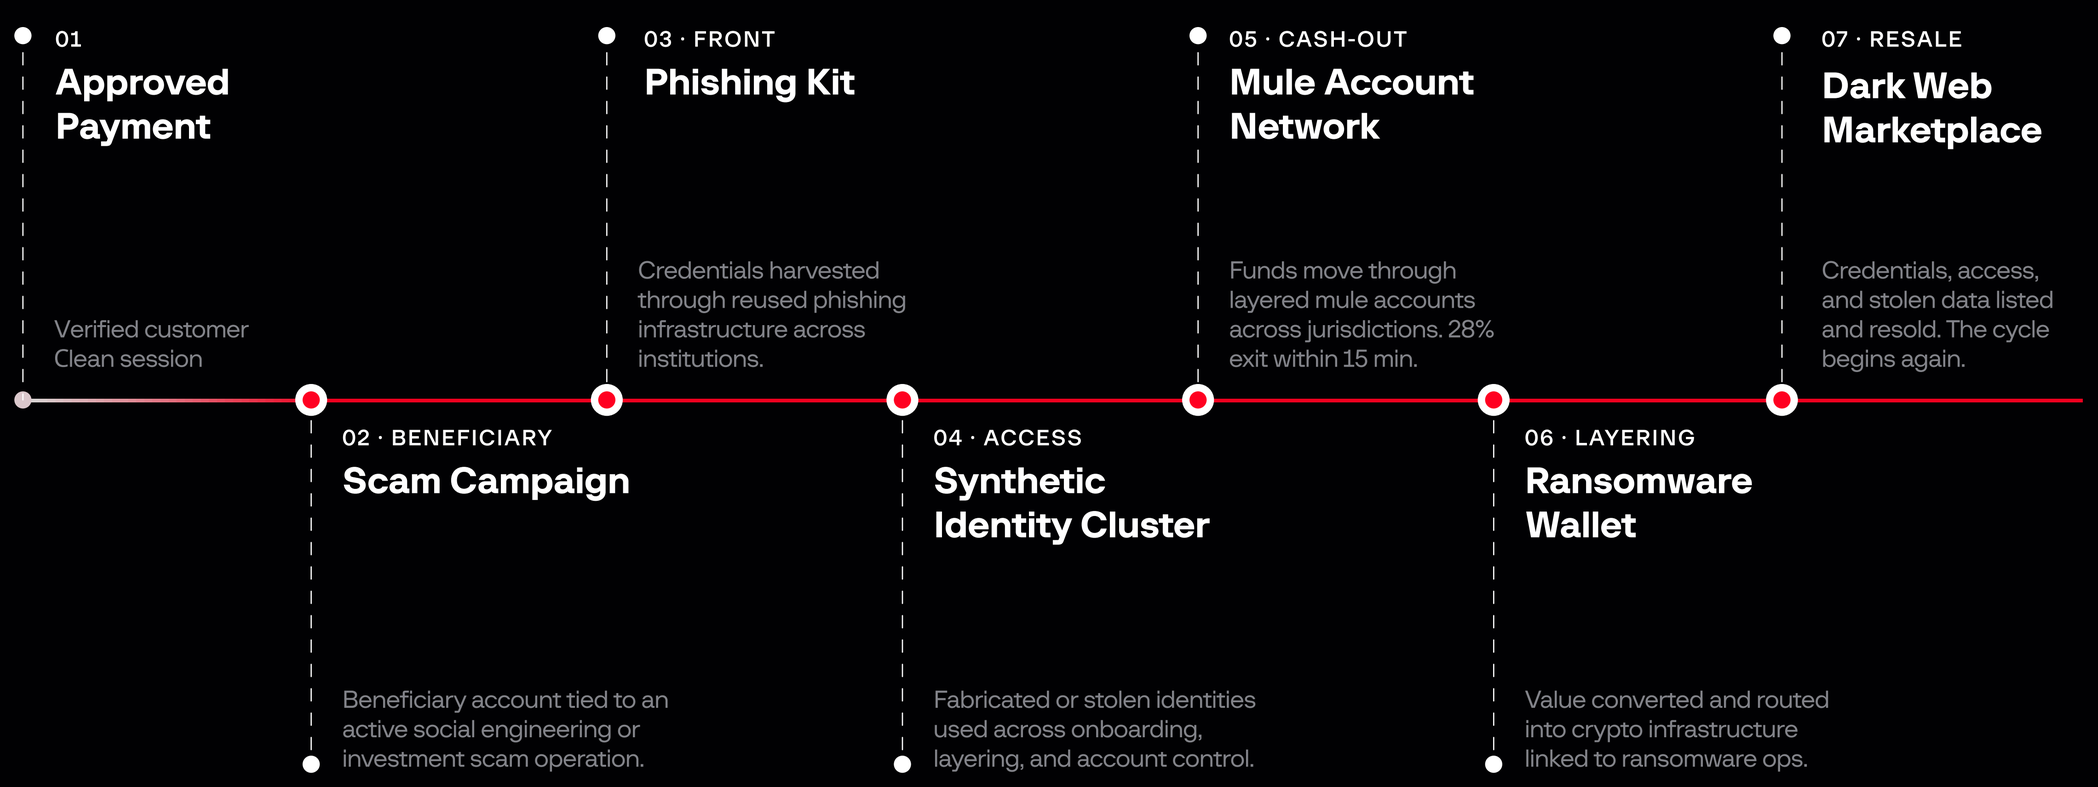
Authorised fraud, where the victim initiates and approves the payment, defeats every transaction-layer control by passing it: genuine customer, genuine device, clean session, correct answers to every challenge. The $579.4 billion lost to bank fraud and scams in 2025- Nasdaq Verafin 2026 Global Financial Crime Report, is dominated by investment, romance, and impersonation scams the victim authorised.

Behind each 'relationship' is the same industrialisation: scripted scam compounds, AI-sustained personas running dozens of simultaneous victims, mule networks weaponised at onboarding weeks before the first payment. 28% of funds are gone within 15 minutes of receipt.

Authorised fraud is routinely cited as the limit of fraud prevention. It is the opposite — the strongest case for Cyber Fraud Fusion, because every decisive signal sits outside the transaction: the beneficiary already flagged across the network, the fake investment platform staged on tracked infrastructure weeks before launch, the grooming arc visible in-session before money moves.

The ecosystem risk chain

The same approved payment is one node in a running criminal supply chain.



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.

The transaction layer asks whether the customer is genuine. In authorised fraud, the only question that matters is whether the destination is.

Every attack type in this section — the conditioned mule account, the malware harvesting session tokens, the synthetic employee requesting elevated access, the scam compound coaching your customer through a transfer — shares one characteristic: the decisive intelligence existed before the transaction. The adversary built an entire operation before your monitoring saw anything. Part II shows exactly what your institution saw, and when.

02 The Gap

Your Architecture Catches Events, Not Campaigns

Three numbers describe the shape of the problem. They are drawn from a representative campaign (a composite of real incidents Group-IB has investigated) not universal averages. But the pattern they show is the architectural default in most institutions today. When read together, they describe a single structural impossibility: the money moves in minutes, but your visibility and your ability to share arrive days later. No amount of effort closes a gap that large. Only a different architecture does.

The Gap	The Number	What It Measures
Internal blindness	**23 days**	From adversary staging infrastructure to your own teams seeing it
Inter-institutional lag	**10–21 days**	From your suspicion to the moment a peer bank can act on it
The vanished window	**15 minutes**	From fraudulent receipt to funds leaving the mule account

The 23-day blind spot

Day	What Happened	Who Saw It	Action taken
Day 01	New mobile-banking trojan circulates via fraudulent app-store listings	Cyber team only	Logged, tagged, filed. No bank infrastructure compromised → no fraud alert raised
Day 12	Malware active on ~3,200 consumer devices. Session tokens harvested, OTP codes intercepted	Nobody	Fraud monitoring sees nothing — no transaction has occurred
Day 23	280 accounts show high-velocity transfers to first-time payees within 36 hours	Fraud team only	Investigated in isolation. Attributed to social engineering. Root cause unknown
Day 31	CERT shares indicator matching Day 01 malware variant	Both — separately	Cyber: "contained." Fraud: "resolved with partial recovery."
Result	**$1.8M confirmed loss**	**Four separate reports**	**Cross-team sessions: ZERO**

Illustrative composite of real Group-IB investigations, not a single measured incident.

The Core Failure: Scoring Fragments vs. Scoring Campaigns

Siloed risk architectures ensure this pipeline remains invisible.

> Cyber teams score intrusion risk. Fraud teams score transaction anomalies. AML teams score structuring patterns. Each scores a fragment. **Nobody scores the campaign.**

A mule account is not created when funds move. It is weaponised weeks earlier at onboarding. New Group-IB research into B2B fintech (datasets based on the France region) reveals a demotivating metric: **1 in 7 business-account sign-ups were assessed as fraudulent**. These operations bypass point-in-time KYC by weaponising a four-stage “clean identity” pipeline that produces synthetic employees and seasoned mules. Because these accounts pair verified personal identities with high-tier business payment capabilities, pre-verified European accounts trade on the dark web for **$300 to $700**.

The Intelligence-Sharing Paradox

The moment fraud becomes shareable, it stops being actionable.

Current fraud reporting frameworks require enough evidence to confirm fraud before intelligence can be shared. But by the time that confirmation exists, the money has already moved and the campaign has progressed. The very process that makes intelligence safe to share often makes it too late to prevent the crime.

A senior AML analyst described the trap precisely: “I file the SAR knowing the funds are already in a mixer. I know it. The regulator knows it. We are both pretending the report means something.”

The traditional reporting timeline: **fraud occurs → 3–7 days to report → 7–14 days to share — 10 to 21 days total**.

The proactive timeline: **suspicious pattern → shared in real time → prevented in real time**.

The time advantage is 10 to 21 days of earlier intervention, and in CFIP deployments, the result is over 90% of those losses prevented.

> **28%**
>
> Of fraudulent funds exit mule accounts within 15 minutes of receipt. By the time confirmed fraud is shareable under current frameworks, the window has been closed for days.
>
> RUSI 2025

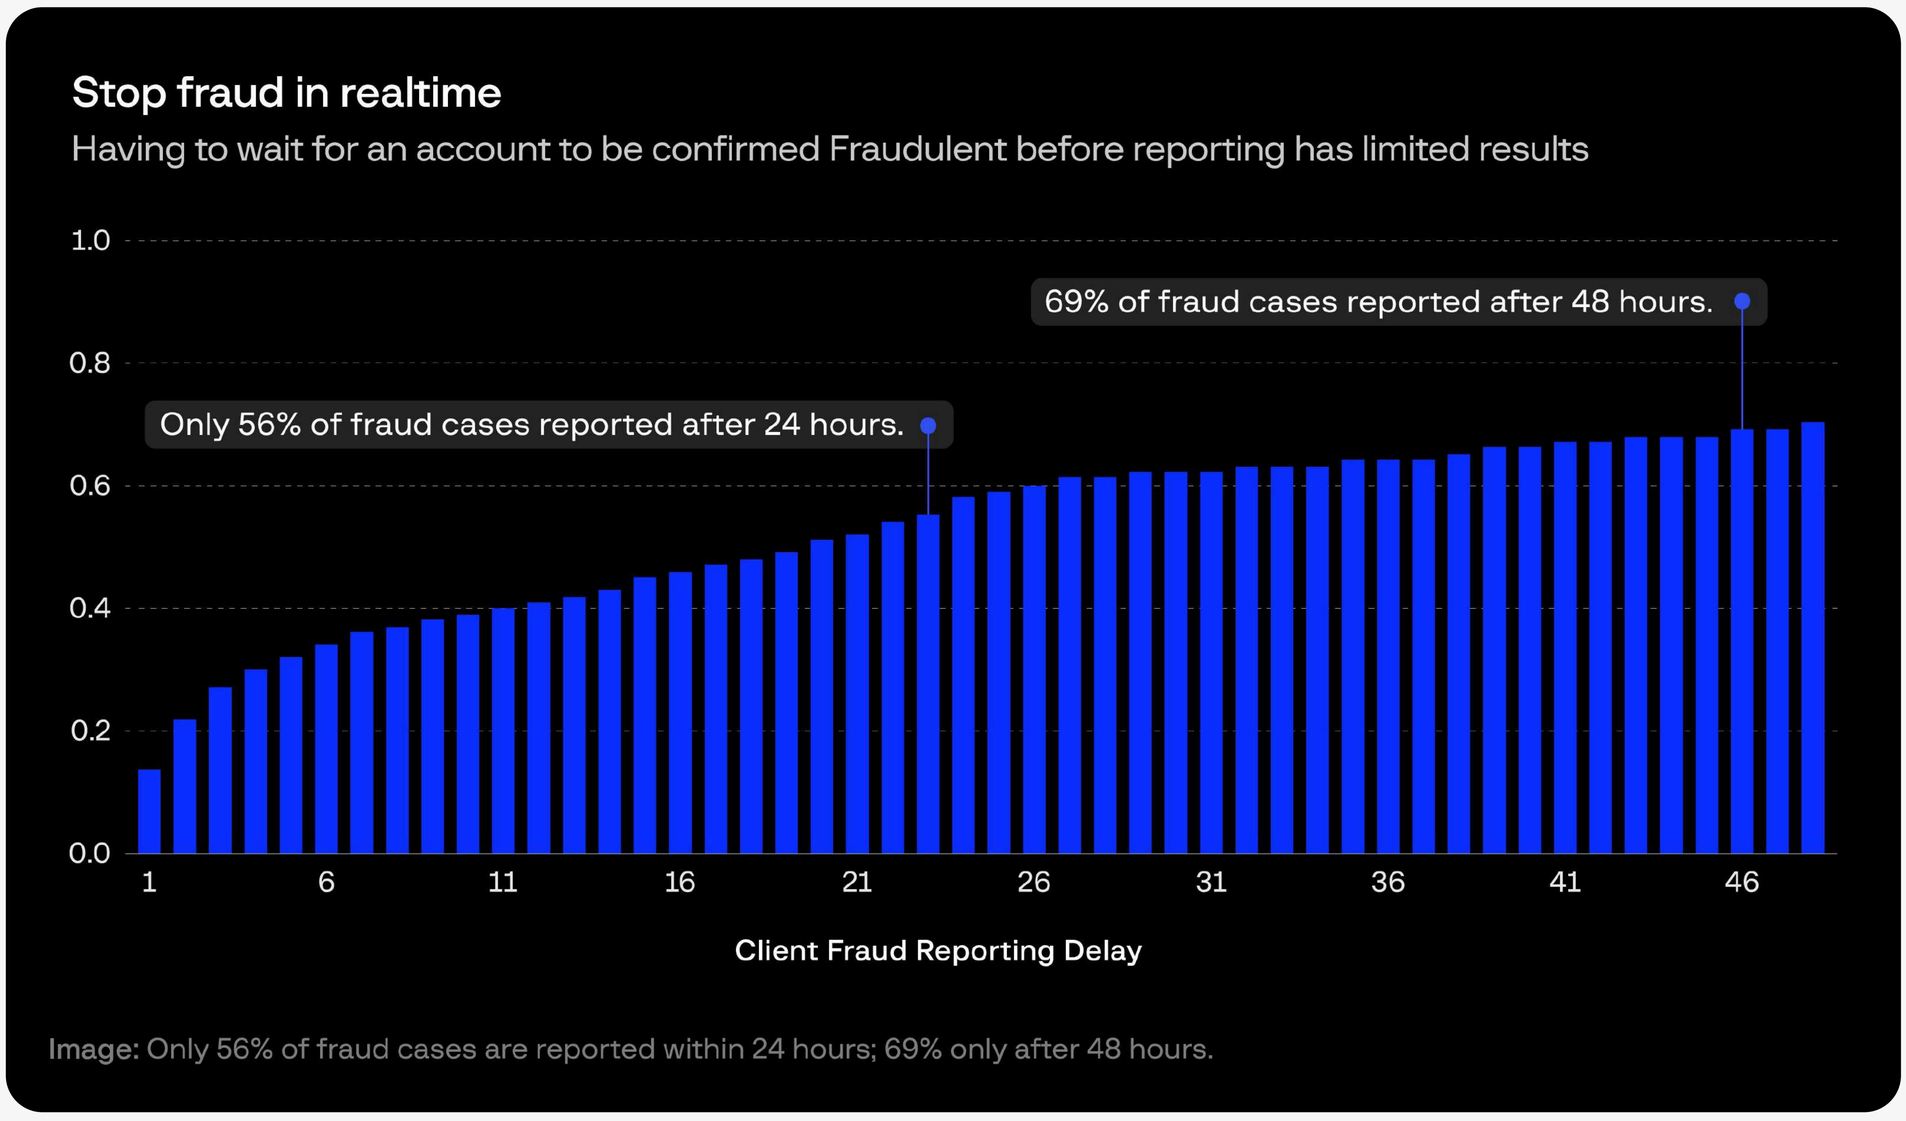


Image: Only 56% of fraud cases are reported within 24 hours; 69% only after 48 hours.

Reporting delays make conventional intelligence sharing ineffective, driving the need for real-time collaborative intelligence models.

Global Approaches to Fraud Intelligence Sharing

GROUP-IB

JURISDICTION	APPROACH	MODEL
UK	Mandatory sharing + Reimbursement	Regulatory mandate
Singapore	Centralized platform (COSMIC)	Government-hosted
Hong Kong	Public-private partnership	Taskforce + IADS
Australia	Economy-wide framework	Multi-sector mandate
USA	Voluntary sharing (FinCEN 314(b))	Safe harbor (liability protect)
EU	Legal framework for data platforms	PSR/PSD3 proposal

Group-IB, 2026

The paradox is real, but it is not unsolvable. Part III shows the cryptographic mechanism that resolves it without institutions exposing any raw customer data.

The decision gap and the vanished window

When **Scattered Spider** targeted financial institutions in 2025, they did not attack banks directly. They attacked the platforms banks trusted — abusing OAuth tokens through Salesforce integrations to access sensitive data across **700+ financial institutions without triggering a single intrusion alert**. The breach lived inside the trusted third-party stack.

Cyber teams see infrastructure being staged. Payment teams see the transaction that arrives after staging is complete. The adversary lives in the space between. The question is no longer 'does this payment look suspicious?' It is: 'does this recipient sit inside an active criminal ecosystem before the funds clear?'

And the window to ask has collapsed. Real-time rails have eliminated the post-initiation review period: SEPA SCT Inst settles in under ten seconds, UK Faster Payments is near-instant, Brazil's Pix processed 16.73 billion transactions in a single month. The risk surface was never a transaction. It is a timeline: a synthetic account opened in March, behaving normally through April, making its first transfer in May, flagged in July, fraudulent the entire time.

The risk decision must be made before payment initiation.

> The 23-day blind spot, the 10-to-21-day sharing lag, the decision gap at authorisation, and the vanished monitoring window are not four problems. They are one problem — an architecture that scores fragments at checkpoints, viewed from four angles. You cannot resource your way out of it. No single point tool fixes it. It requires a different operating model.

[EVALUATE]: The Kill-Chain Coverage Test For Your Existing Tools

The gaps above are not four separate problems, but are architectural problems viewed from four angles. Before you evaluate new solutions, you need to know where your current stack actually sits.

The most revealing assessment is mapping your controls against the MITRE Fight Fraud Framework (F3) — the standardized taxonomy of financial fraud TTPs launched in April 2026. F3 organises fraud into seven phases: Reconnaissance, Resource Development, Initial Access, Defense Evasion, Positioning, Execution, and Monetization.

The critical question: Which of these seven phases does each of your tools actually see?

If your tools collectively leave Positioning (adversary preparing inside your environment) and early Monetization (funds routing before extraction) uncovered, you have the exact decision gap described above.

Which fraud phases does each part of your stack actually see?

Your Current Tool	What It Sees	What It Misses	The Gap
WHAT MOST INSTITUTIONS RUN TODAY			
Enterprise Fraud Management (EFM)	Transaction and account anomalies at authorisation	Reconnaissance → Positioning	Campaign invisible until it reaches the payment
KYC platform	Account opening	Resource Development	Synthetic-identity pipeline invisible
AML	Structuring patterns	Pre-money-movement staging	Mule networks weaponised at onboarding
WHAT ACTUALLY CLOSES THE GAP — THE FUSION LAYERS			
Threat Intelligence + Digital Risk Protection	Resource Development / Reconnaissance — staged phishing kits, harvested credentials, adversary infrastructure	This is the pre-attack layer	Sees the campaign forming weeks before the transaction
Fraud Protection	Positioning / Execution — device intelligence, behavioural biometrics, scam-call and RAT signals in-session	This is the in-session layer	Catches the attack at the session, before the payment authorises

The test isn't which tools you own — it's whether, together, they see every phase. Most institutions cover Execution and Monetisation, and leave Reconnaissance, Resource Development, and Positioning uncovered. That is the decision gap.

Four Structural Criteria
Judge your stack on these dimensions:

A. Temporal Reach — When do you see it?

Stage 1 (Reactive): After transaction completes — reimbursement/loss recovery.

Stage 2 (Detective): At transaction authorisation — rules engines, scoring.

Stage 3 (Predictive): During positioning/preparation — behavioural anomalies, device intelligence.

Stage 4 (Preemptive): During resource development/reconnaissance — threat intelligence on adversary infrastructure.

Most institutions are Stage 1–2 on fraud, Stage 3–4 on cyber, and never the twain shall meet.

B. Cross-Domain Signal Flow — Does intelligence move?

Can a cyber signal, such as, a device flagged by threat intelligence, malware or a RAT detected in-session automatically raise a fraud-risk flag before the payment, without a manual handoff?

Can your fraud platform consume threat intelligence feeds about mule account patterns?

Do your KYC, fraud, cyber, and AML teams share a common incident taxonomy?

If the answer to any is "no," you have a silo problem, not a tool problem

C. Campaign Continuity — Can you trace across time?

Can you link an onboarding session in March to a payment pattern in June?

Can you connect a device fingerprint from login to the same fingerprint at account opening?

Can you maintain identity continuity across dormancy periods?

If your tools reset their analysis at each session, you cannot see campaigns.

D. Ecosystem Visibility — Do you see beyond your perimeter?

Can you query whether a beneficiary account has been flagged by other institutions?

Check if a payment destination sits on known criminal infrastructure?

Receive pre-confirmation intelligence without exposing customer PII.

03

The Operating Model

What cyber-fraud fusion would have done

The institutions furthest ahead share one architectural characteristic: they stopped treating fraud, cyber, AML, and KYC as separate functions that share information periodically, and started building a single intelligence layer that all four draw from continuously. Here is what that produces, measured directly against the gaps Part II established:

Adversary Stage	The Event	Siloed Response	Fusion Response
Pre-attack	**Day 01** — Malware kit circulates	Cyber logs it. Fraud never told.	Threat Intelligence detects kit on underground markets before deployment. The fraud team primed with device indicators.
Pre-attack	**Day 12** — Malware active on 3,200 devices	Nobody. No transaction yet.	Device fingerprints from infected sessions flagged automatically. Watch-list live.
In-session	**Day 23** — 280 accounts show high-velocity transfers	Fraud investigates in isolation	Fraud Protection blocks sessions in real time. Malware family identified. Upstream Threat Intelligence updated.
Cross-institution	**Day 23+** — Funds reach mule accounts	SAR filed 3–7 days later. Funds gone.	Cyber Fraud Intelligence Platform flags receiving tokens instantly across the network. Mule rings surfaced. Funds frozen.
Result	**Four teams. $1.8M lost. Zero cross-team sessions.**		**One loop. > 90% of losses prevented. Campaign closed.**

Four teams, four moments, one signal carried across the whole architecture without a manual handoff or a 23-day gap. The credential is not just blocked at your door; the operator behind it is taken off the board for everyone.

LAYER 1

Pre-Attack

THREAT INTELLIGENCE

DIGITAL RISK PROTECTION

Detect the attack while the adversary is still preparing

The capability: Group-IB Threat Intelligence (TI) + Digital Risk Protection (DRP)

It is the upstream sensor layer that watches the adversary while they still prepare. Threat Intelligence draws on dark-web forums, criminal Telegram and Discord channels, closed marketplaces, ISP-level sensors, malware sandboxes, and operational collaboration with law enforcement. **Digital Risk Protection detects phishing pages, fake mobile apps, and brand abuse, then routes them to takedown workflows delivering an average 85% pre-trial takedown rate and eliminating more than 20,000 violations every day for over 450 protected brands.**

In the $1.8M campaign

The malware kit was already in the Group-IB threat database on Day 01. With the pre-attack layer active, the fraud team would have received a pre-positioned watch-list of infected device fingerprints on Day 01, not Day 23. The 23-day gap closes to minutes.

The strategic effect is adversary economics: every phishing kit taken down before it earns, every mule flagged before activation, raises the attacker's cost-per-successful-fraud. That is the only metric the criminal economy ultimately responds to.

LAYER 2

In-Session

FRAUD PROTECTION

Intercept fraud at the session, before the payment authorises

The capability: Group-IB Fraud Protection (FP)

When prevention fails, and at scale it always fails partially, the next checkpoint is the customer's own session. Fraud Protection fingerprints devices, monitors OS and sensor configuration, detects malware, bots, and Remote Access Trojans in real time, and performs anonymised behavioural analysis (mouse, keyboard, touchscreen) to identify whether the human in the session is the real customer, a fraudster, or an automated framework. Scam-call identification detects when a customer is being coached through a fraudulent transfer by a third party — the signature of authorised-push-payment fraud. It protects more than 300 million users worldwide.

In the $1.8M campaign

Device fingerprints from the 3,200 infected sessions would have been flagged against the Day 01 TI indicators. FP blocks the anomalous sessions on Day 12 — before a single transfer fires on Day 23. Every blocked session also feeds upstream: the malware family is identified and the TI database updated for the next institution.

Every blocked bot session, every detected RAT becomes a new signal, fed upstream into the pre-attack layer and downstream into the cross-institution layer. The session is not just where you stop fraud. It is where fraud tells you what it is doing, if you have the architecture to listen.

Independent validation (Forrester Total Economic Impact™ study): 10–20% more fraud attempts blocked, 20% fewer false positives, 30% fewer one-time passwords required, 130% ROI within six months.

LAYER 3

Cross-Institution

CYBER FRAUD INTELLIGENCE PLATFORM

Share suspicions across institutions, before confirmation and not after

The capability: Cyber Fraud Intelligence Platform (CFIP)

By the third stage, most institutional defenses end. The transaction has cleared; funds are in the mule account, and 28% of the time they are gone within fifteen minutes. A consortium model that trades in confirmed-loss reports cannot meet that clock. CFIP is the network layer of fusion: a privacy-preserving exchange that lets participating banks, payment providers, telcos, and government identity services share suspicions in real time, without ever exposing raw customer data.

In the $1.8M campaign

The receiving mule accounts including the one that opened this paper — would have been queried against the CFIP network at the $12 test, the week before the strike, not at Day 23. The same operator's footprint across multiple banks is unmistakable in aggregate. The mule ring would have surfaced before the first $47,000 transfer was ever authorised.

How the privacy paradox is resolved: Not with policy, but cryptography

CFIP's core protocol (patented Distributed Tokenisation, independently audited by Bureau Veritas against GDPR Article 26) irreversibly tokenises sensitive identifiers inside each participant's environment before any signal leaves the building. No single participant holds the full salt. No third party can re-identify the token. The PII never leaves the originating institution. The actionable intelligence travels across the ecosystem in minutes.

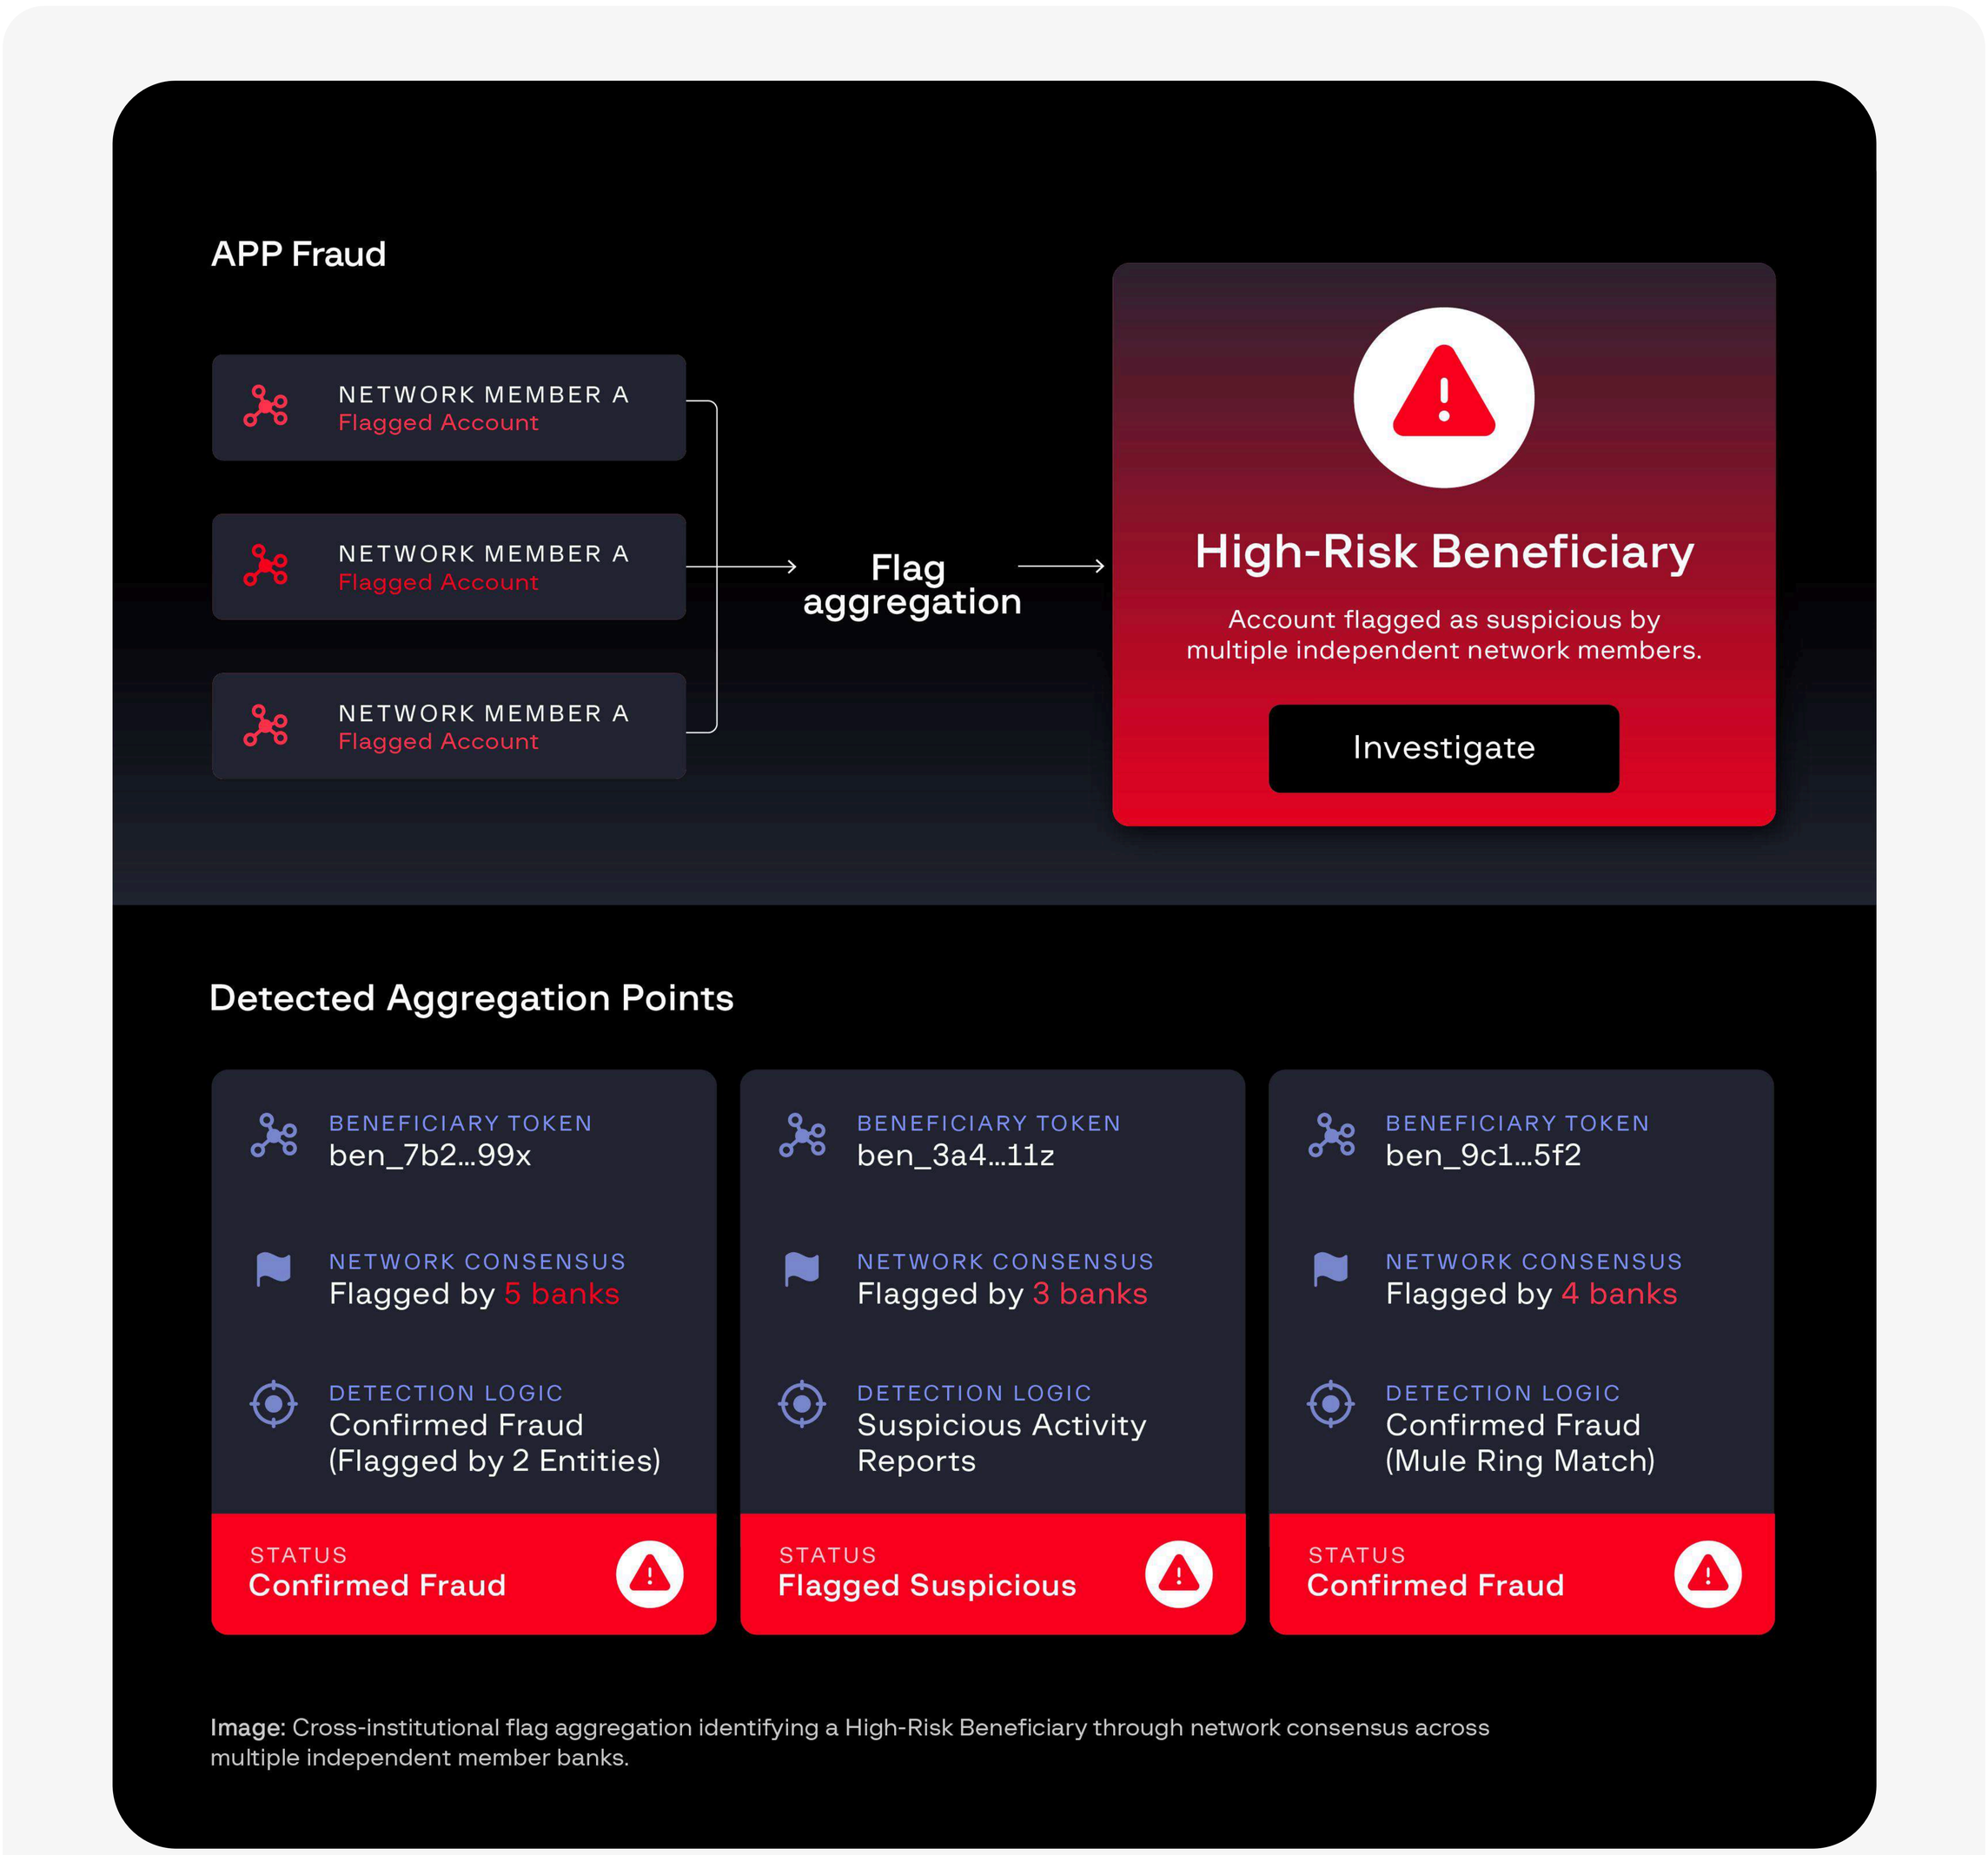


Image: Cross-institutional flag aggregation identifying a High-Risk Beneficiary through network consensus across multiple independent member banks.

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

The gap it closes: The intelligence-sharing paradox. What CFIP produces that no single institution can build alone includes:

- **Mule warm-up detection:** The conditioning that looks legitimate inside one bank is unmistakable across the network.
- **Cross-institution beneficiary scoring**: When several institutions flag the same recipient token in a short window, it surfaces as a confirmed mule-ring match.
- **Cyber-enriched signal fusion:** Each query enriched against the Threat Intelligence database — compromised cards, infostealer credentials, devices tied to known operators.
- **Cross-sector identity verification:** A 'new customer' revealed as a pre-paid SIM activated thirty days ago, already linked to a mule account elsewhere.

Pilot evidence: A national deployment runs **180,000 checks per day**, surfaces **300–400 known mule cards per day**, and is projected to prevent **$10–15 million annually** with only 5% of network members actively checking. Production latency across 34 banks and 12 payment systems: mean 22–38 milliseconds, p95 under 42 milliseconds.

A Common Language: The MITRE Fight Fraud Framework™

An adversary-centric operating model needs more than shared technology; it needs a standardised operational language. **Group-IB serves as an initial data contributor to the newly released MITRE Fight Fraud Framework™ (F3)**, developed by the MITRE Corporation. By integrating the F3 taxonomy directly into the Fraud Matrix, institutions can systematically map real-world investigation data to standardised fraud techniques. This bridges the historic communications gap between security operations centres and fraud-risk units.

When a new technique is classified via the framework, the platform automatically pairs it with live adversary infrastructure tracking and exact technical counter-measures, letting the global financial ecosystem classify, predict, and neutralise adversary behaviour long before a payment is initiated.

04 Proof

The 2025 Operational Record

In 2025 alone, Group-IB's operations, through 11+ Digital Crime Resistance Centres across 60 countries, in direct partnership with INTERPOL, Europol, and national agencies, produced the following. These were not investigations that closed after the loss was booked. They disrupted live criminal operations, seizing infrastructure and arresting operators before campaigns could scale to their next wave of victims.

Metric	2025 Result
Cybercriminals arrested	1,809
Malicious infrastructure nodes dismantled	34,838
Estimated financial losses prevented	US$100M+
Victims confirmed protected by law enforcement	310,643

Case Study: A National Fraud-Defence Hub

A central bank in a major Central Asian country faced a fragmented defence against coordinated financial crime — data silos between institutions creating blind spots, the reactive model meaning that by the time fraud was discovered, the money was gone. The bank deployed CFIP to establish a national, real-time intelligence hub using Distributed Tokenisation to share intelligence on suspected and confirmed mule accounts.

Even in pilot, with only 5% of member banks performing real-time checks, the platform delivered:

- A 5× increase in fraud reports shared across the network within six months.
- 300–400 mule accounts identified daily, across 180,000 checks per day.
- $10–15M in losses prevented annually at pilot scale.

At full adoption, the projection reaches $100–300M annually.

05

What this changes

From a posture you defend to one you can prove

The fraud disruption maturity model.

Fusion is a capability you mature into, not a team you stand up. The four-stage model below lets a Head of Fraud and a CISO agree on where they actually are today, and makes the next move a small, fundable step rather than a leap to an end-state no one has budget for.

Stage	Defends	Acts	Primary metric
1 · Reactive	Events	After the loss	Confirmed losses
2 · Proactive	Transactions	At the payment	Detection rate
3 · Predictive	Campaigns	During staging	Cross-domain MTTD
4 · Disruptive	The ecosystem	Before launch	Disruption rate

Your stage is set by your worst silo, not your best tool. Pilots do not count; defaults do. It is not a converged institution until converged intelligence is the standard input to a payment decision.

The budget conversation no one starts early enough

There is a quieter argument the buying committee raises late, and it belongs to the CFO and the CRO: cost predictability. Three teams, three tools, and three procurement cycles produce a sprawl of per-seat, per-query, per-feed line items — and a budget you keep re-explaining. A consolidated platform — one data lake covering threat intelligence, fraud, detection, and response, with unlimited users, APIs, hunting rules, and takedowns — turns that sprawl into a single, forecastable number.

How the privacy paradox is resolved: Not with policy, but cryptography

This produces a structural cost advantage before it prevents a single fraud loss. Institutions that buy against the architecture stop paying twice for the same dark-web feed, the same brand monitoring, the same enrichment source — and end up on a fundamentally different cost base from those still assembling point solutions. **Gartner noted this directly in naming Group-IB a Leader in the 2026 Magic Quadrant™ for Cyberthreat Intelligence Technologies — citing the breadth and cost predictability of the unified platform, the Cyber Fraud Fusion model by name, and a roadmap shift 'from detection to prediction-first defence.'**

The accountability shift is already here

Regulators in every major market are arriving at the same conclusion, independently. Not fines for what institutions did wrong — fines for what they failed to build.

Jurisdiction	Framework	Core obligation
UK	Payment Systems Regulator (PSR)	Mandatory sharing + direct reimbursement liability — PSPs that fail to prevent bear the loss
EU	PSR / PSD3	Legal framework for cross-institutional data sharing; continuous session-wide monitoring mandated
Singapore	Shared Responsibility Framework (SRF)	Bank + telecom joint liability — shared accountability regardless of where the gap occurred
Australia	Scams Prevention Framework (SPF)	Cross-sector penalties — banks, telecoms, and platforms all liable; failure-to-prevent triggers fines
USA	FinCEN 314(b) + state mandates	Voluntary sharing with safe harbour; failure-to-prevent pressure growing at state level
UAE / KSA	CBUAE / SAMA frameworks	Real-time monitoring obligations — cross-institution reporting now a baseline expectation

How the privacy paradox is resolved: Not with policy, but cryptography

The fusion model is an intelligence layer above existing systems, not a substitute for them. Your rules engine, transaction-monitoring platform, and case-management tooling remain in place; they become consumers of converged intelligence rather than isolated producers of fragments. Adoption requires three things.

Data readiness: Telemetry, not just transactions — device and session signals, identity and onboarding data, payment context, and cyber telemetry, available at decision time. Most institutions already generate all four; the work is connecting them.

Clear ownership: A single accountable executive for cross-domain financial-crime intelligence and a fusion working group with the authority to act on converged signals. This does not mean merging departments; it means one escalation path and one entity-level case view.

A legal basis: Established once, per jurisdiction, at onboarding. Cross-institutional signal sharing through Distributed Tokenisation operates within data-protection and financial-crime confidentiality regimes — tokens, not PII, travel. Group-IB provides the Bureau Veritas validation documentation and jurisdiction-specific legal mapping to support the review.

The First Move

How can you start the fusion journey this quarter with the team you already have?

The "Virtual Fusion Team" is a bridge to validate the data loop, not a permanent operational state. Group-IB's goal for the 90-day ramp is to establish **automated playbooks**. If a customer device triggers a behavioral biometric anomaly (e.g., typing rhythm shifts dramatically, indicating an active scam call or RAT), the system handles the cross-silo logic programmatically. It updates the central profile, alerts fraud ops at login, and readies the token for the peer network. The humans only step in at the very end of the loop.

The virtual fusion team

Existing role	What they bring	Typical time
Threat-intelligence analyst	Pre-attack signal: attacker infrastructure, leaked credentials, dark-web chatter, mule indicators	**½ day / week**
Fraud operations analyst	In-attack reality: live typologies, rule performance, false-positive pain, what is actually hitting customers	**1 day / week**
SOC / detection engineer	Telemetry, device and session signal, and the plumbing to route data between functions	**½ day / week**
AML / financial-crime analyst	Post-attack view: mule behaviour, money movement, the regulatory frame for any sharing	**½ day / week**
Digital-channel product owner	The customer-experience constraint: where friction is acceptable and where it is not	**As needed**
Fusion lead (rotating)	A single named owner of the ritual and the scorecard — a hat worn by an existing manager, not a new hire	**2–3 hrs / week**

Governance at this stage is one meeting and one artefact, not a charter and a steering committee. Over-formalising is the most common way these efforts die. A standing fusion review of 30 to 45 minutes, weekly or fortnightly. One named owner. A single shared workspace where watch-lists, indicators, and decisions live where both teams can see them.

The 90-day on-ramp

Days 1–30

Map the silo

Inventory every dark-web, brand-monitoring, threat-intelligence, and fraud-detection tool across the institution; identify where two teams pay for substantially the same capability. Stand up the joint fusion working group with one named owner. Make MITRE F3 its standing vocabulary. Pull one recent account-takeover incident through the F3 taxonomy end-to-end, marking every point where intelligence existed but did not reach the team that needed it.

Days 31–60

Run one closed loop in production

Choose one fraud typology (authorised-push-payment, account-takeover, or mule onboarding) and one channel. Wire pre-attack intelligence, in-session detection, and post-attack mule signals into a single playbook. Measure detection lift, false-positive rate, and time-to-block against the baseline. Begin participating in at least one cross-institution intelligence network.

Days 61–90

Prove the economics, scale the architecture

Quantify avoided-loss, false-positive, and friction impact against a control — bring the numbers to the risk committee as a recommendation, not a request. Commission an F3-based maturity assessment. Adopt the procurement principle: net-new capability is bought against the fusion architecture, not the silo.

> Ninety days is enough to prove the model on one typology and one channel. It is not enough to finish the transition — but it is enough to make the case to finish it. The capability compounds: every shared signal makes the next one cheaper to act on, and every closed loop makes the next one easier to wire.

The Decision

Three Questions for Your Leadership Team

Each is designed to be answered honestly in a single meeting. If the answer to any is 'no,' the architecture is the reason.

1. Does your fraud, AML, cyber, and KYC intelligence converge at the entity level, in real time? **Not in quarterly reports, but at the moment a payment is being decided.** If the answer is no, your architecture is scoring fragments of a campaign it cannot see in full.

2. Do you have visibility into how your institution appears from the adversary's side? Which customer credentials are on underground markets today? Is your brand being staged for an upcoming campaign? Without adversary-side visibility, the adversary has better intelligence about your customers than you do.

3. Can you act on a suspicious signal before it becomes a confirmed loss? **Not file a report, act: block an account in warm-up, stop a payment to a beneficiary already flagged across three other institutions, identify a synthetic identity before it activates.** If your intelligence cycle starts with confirmed fraud, the window to prevent it has already closed.

Current state

- Fraud team detects events
- AML reports after funds move
- Cyber intelligence never reaches fraud ops
- KYC is a one-time checkpoint
- Compliance is a retrospective audit

Fraud loss is a cost of doing business

Intelligence-led state

- Cyber-fraud fusion sees campaigns
- Suspicious signals shared pre-confirmation
- One intelligence layer feeds all four functions
- Identity trust is continuously re-evaluated
- Monitoring generates the evidence base in real time

Fraud disruption is a competitive advantage

The question is no longer whether intelligence-led prevention is possible. It is whether your architecture, your tooling, your team structures, and your vendor agreements are built to support it. Under PSR and its global counterparts, the cost of that answer being "no" now lands directly on your balance sheet.

Start Here

Loss reduction. Regulatory-exposure reduction. Competitive differentiation. All three unlock from the same first move.

Move toward Cyber Fraud Fusion as an operating model: fraud, cyber, AML, and KYC as one intelligence layer, grown into one step at a time, not through a rip-and-replace. Start with any of the three below.

Threat Exposure Assessment

Understand exactly how your institution appears from the adversary's 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 most exposed right now. Delivered by Group-IB analysts with direct access to criminal-ecosystem intelligence.
https://www.group-ib.com/contacts/

Cyber Fraud Intelligence Platform (CFIP) and Fraud Matrix Demonstration

See what real-time, privacy-preserving cross-institutional signal sharing looks like in practice, and what it would have flagged in your last three major fraud incidents. CFIP is the only GDPR-validated solution that enables pre-confirmation signal sharing without exposing PII. The Fraud Matrix maps the full adversary kill chain so your teams converge on campaigns, not fragments.

Fraud Exposure Calculator

Two inputs. No benchmarks. Enter your institution's own operating numbers and see the cost of your exposure under the pre-crime intelligence model, the gap between what your books record and what your data already knows, before your next budget conversation.
https://www.group-ib.com/fraud-calculator/

> The institutions that survive this decade will be the ones that see campaigns, not events, and act before the transaction, not after. That is the operating model. The first step is yours — and it does not require permission.

TECHNICAL ANNEX

FOR ENTERPRISE ARCHITECTS

Deploying Cyber Fraud Fusion Alongside Your Core

A reference architecture for adding intelligence-led fraud prevention without a core banking overhaul.

Companion to Before Fraud Transacts: Enabling Proactive Prevention in the Age of Industrialised AI

The premise. Adopting the fusion model does not require replacing your core banking platform. Group-IB Fraud Protection (FP) and the Cyber Fraud Intelligence Platform (CFIP) deploy as an intelligence layer that runs parallel to your existing systems. Your transaction-monitoring engine, rules engine, and case-management tooling remain in place and become consumers of an enriched risk signal — they are not altered, and they retain the final authorisation decision.

Reference data flow

User session
Web / mobile / 3-D Secure

→

Client-side module
Web Snippet or Mobile SDK

→

Preventive Proxy
At gateway / load balancer

↓ device + behavioural telemetry

↓ async, parallel to transaction

↓ risk score on session token

FP Processing Hub
Cloud, on-prem, or hybrid

→

CFIP network query
Tokenised, < 100 ms

→

Your existing risk engine
Final auth / freeze decision

How to read this. Telemetry is collected at the presentation layer and processed in parallel with — never in the critical path of — transaction authorisation. The enriched risk context is appended to the session token at the gateway. Your existing risk engine consumes that score and makes the final decision. Nothing in the core banking environment is modified.

The three components

Component	What it is	Where it sits
Client-side module	A lightweight Web Snippet (web portals) or Mobile SDK (apps) embedded in the presentation layer. Integration does not require changing application logic. Not agentless — it is a client-side module you add to your own code.	Web portal, mobile app, or 3-D Secure flow
Preventive Proxy	Intercepts session, device, and behavioural telemetry before requests reach the core. Deploys as a loop on the load balancer, or in the gap between the load balancer and the application backend.	Application gateway / load balancer
Processing Hub	Generates the verdict and risk score from collected telemetry. Available as cloud (SaaS), on-premises (standalone), or hybrid — e.g. Preventive Proxy on-prem, Processing Hub in the Group-IB cloud.	Per your deployment model

Three deployment models — choose by data-residency need

Model	Description	Typical fit
Cloud (SaaS)	Fast integration with cloud infrastructure in the customer's country. Group-IB-hosted processing.	Speed to pilot; no on-prem footprint
On-premises	Standalone deployment fully inside the customer's infrastructure. Monitoring delivered by Group-IB.	Strict data-residency or air-gap requirements
Hybrid	Customised split — e.g. Preventive Proxy on-premises, Processing Hub in the Group-IB cloud.	Balance of control and operational simplicity

What this means for your stack

No core replacement	FP and CFIP integrate with any risk-management platform that exposes an API, in both pull and push modes. CFIP's microservice architecture connects to case-management, risk-engine, and transaction-monitoring tools without rebuilding infrastructure.
Your engine keeps the decision	The fusion layer supplies an enriched risk signal; your existing rules engine makes the final authorise-or-freeze call. Peer signals from CFIP enter as a corroborating input, never as the sole determinant.
Runs parallel, not in-line	Telemetry collection and network queries run asynchronously alongside transaction processing, so the fraud layer does not add latency to the payment path itself.
The < 100 ms figure is CFIP's	The sub-100-millisecond latency refers specifically to CFIP's cross-institutional tokenised query — the privacy-preserving network lookup — not the in-session FP telemetry path. Quote it against the network layer, not the edge.

Group-IB HTCT 2026 · Group-IB Proactive Fraud Prevention Whitepaper 2026 CFIP Solution Brief · GASA Global State of Scams 2025 · RUSI 2025 · Forrester Total Economic Impact™ · UNODC / World Bank StAR Initiative · MITRE Fight Fraud Framework™ (F3), April 2026, Gartner does not endorse any vendor, product, or service depicted in its research publications and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. GARTNER and Magic Quadrant are registered trademarks and service marks of Gartner, Inc. and/or its affiliates in the U.S. and internationally and are used herein with permission. All rights reserved.

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