Fortify.

THE CONVERGENCE PLAYBOOK

A best-practices companion guide

How banks, payment processors, remittance firms and crypto businesses can apply the learnings from the white paper.

A companion to the white paper

"Stronger Together: What Fraud and Cyber Can Learn from Each Other", by Ken Palla.

Published with Fortify · 2026

wearefortify.ai

Fortify.

CONTENTS

Inside this guide

01 How to use this guide

02 Banks

03 Payment processors

04 Remittance firms

05 Crypto and digital-asset firms

06 From shared intelligence to shared action

This paper is product-agnostic. It is written for practitioners on both sides of the fraud and cyber divide, and every framework referenced in it is publicly available.

SECTION 01

How to use this guide

The white paper made a simple case. Fraud, cyber and AML teams are fighting the same people, and they win when they work as one. This guide turns that case into practice. Each chapter takes a type of institution, narrates one attack the way the criminals run it, shows how it fragments into separate alerts for separate teams, and sets out a workflow, the signals to share, and how to tell it is working.

THREE IDEAS RUN THROUGH ALL OF IT.

01 **The criminals don't respect your org chart.**

The network that socially engineers a customer, drains an account and launders the proceeds is often one operation. Three teams each see a fragment of it.

02 **The data needs to be shared.**

These teams have struggled to share for a practical reason: their data sat in systems that couldn't talk to each other. A shared data estate changes what is possible.

03 **AI has to explain itself.**

AI now finds weaknesses in minutes, for attackers and defenders alike. The teams that come out ahead are the ones whose automated decisions can be explained and traced.

SECTION 02

Banks

The threat picture

Banks carry the widest threat surface of any institution in this guide. Deepfakes are testing account-opening controls. Authorised push payment scams move customer money out in seconds. Business email compromise still costs commercial clients six and seven figures. Underneath all of it sit AML obligations that never pause: transaction monitoring, sanctions screening, suspicious activity reporting. The teams handling these usually sit apart, on different systems, with different incentives. The cost of that separation is easiest to see in a single case.

Anatomy of an attack: the scam that looks like three incidents

A retail customer takes a call from someone claiming to be the bank's fraud team. The number is spoofed, the script is rehearsed, and over forty minutes she is talked into moving her savings to a “safe account” to protect them. She authorises the faster payment herself, so every individual control she passes through reads as legitimate.

The money lands in an account opened three weeks earlier with a synthetic identity that cleared onboarding because an injected deepfake selfie beat the liveness check. Within minutes it is layered across four accounts and out through a crypto on-ramp. Separately, the same group has been credential-stuffing the bank's mobile API for a fortnight, and one of the mule accounts was opened from a device the security team had already logged in an account-takeover cluster.

To the bank, this is three events. Fraud opens a scam case for the victim. Days later AML's monitoring flags the mule account's rapid layering as its own alert, with nothing tying it to the scam. Cyber saw the credential stuffing and the repeat device and filed it as noise. The device that links the mule account to the account-takeover cluster is the single thread running through all of it, and no one pulls it, because no one team can see both ends.

What each team sees, and its blind spot

Team	What it sees	Its blind spot
Fraud	The victim, the scam narrative, the outbound payment.	Where the money went, and who owns the receiving account.
Cyber	The repeat device, the credential stuffing, the API abuse.	That the device links to a funded mule account and a live scam.
AML	The mule account's rapid layering and the crypto off-ramp.	That the inbound funds came from a scam victim, not real activity.

The joined-up workflow

A practical sequence to run the moment a scam case opens, so the chain is read as one case rather than three.

	Owner	Action	Passes to the others
01	**Fraud**	Open the scam case on the outbound payment. Capture the receiving account, sort code and any Confirmation of Payee mismatch.	Receiving account and CoP mismatch to AML.
02	**AML**	Check the receiving account against open mule indicators and live layering alerts. Report whether it is a known or suspected mule and where the funds moved.	Mule status and layering path to Fraud and Cyber.
03	**Cyber**	Run the receiving account's onboarding and login devices against account-takeover and bot clusters.	Linked devices, IPs and accounts to Fraud and AML.
04	**All three**	Decision point. If device or account links exist, treat it as a network case, not a lone scam. Freeze the linked mule accounts and file one SAR covering the whole chain.	A single joined case record.
05	**All three**	Feed the confirmed indicators back into each team's detection, so the next attempt is caught earlier.	Updated rules and blocklists.

Signals worth sharing

- → **Fraud to AML:** receiving account and sort code, Confirmation of Payee mismatch, scam typology, the beneficiary the victim named.
- → **Cyber to Fraud and AML:** device fingerprint, IP and ASN, the login-anomaly cluster ID, credential-stuffing hits on the account.
- → **AML to Fraud and Cyber:** mule confirmation, the layering path, the cash-out or crypto endpoint, linked accounts.

How you'll know it's working

01 Time from scam report to freezing the receiving account, measured in hours rather than days.

02 The share of scam cases that get linked to an identified receiving account.

03 One SAR per criminal chain, instead of scattered, disconnected alerts.

04 Repeat devices and accounts caught at onboarding, rather than after the money has moved.

Where to start

01 Map who owns each stage of the scam-to-mule-to-laundering chain, and where the handoffs currently drop.

02 Run the workflow above on one real recent scam case as a tabletop exercise.

03 Agree the three signal lists above, and open the data both ways.

04 Add “show your working” to the criteria for any AI decisioning tool.

SECTION 03

Payment processors

The threat picture

Processors live on volume and thin margins, which makes manual review the enemy of the business model. The risk sits at two points: who you let onto the platform, and what flows across it once they are on. Merchant onboarding fraud and transaction laundering run alongside card testing, credential stuffing and API abuse at scale. The AML duties on the merchant base don't lift because the volume is high. As with a bank, the cost of separate teams shows up in a single merchant.

Anatomy of an attack: the merchant that is really three problems

A merchant onboards through the self-serve flow with clean-looking documents, a real but stolen company identity, and a card-acceptance request that clears KYB. For two weeks it trades modestly and builds a normal-looking history.

Then volume spikes overnight: thousands of small card payments, many of which fail. It is card testing, run through the merchant to validate stolen card numbers. The cards that work are then charged for “sales” that are really laundering, washed through a legitimate-looking storefront. The traffic hitting the payment API during the spike comes from the same bot infrastructure the security team has been watching abuse a different endpoint.

To the processor this is, again, three events. The risk team sees a merchant with rising chargebacks. Transaction monitoring flags unusual settlement as a possible laundering case. Security sees API abuse and rate-limits it. The merchant that ties all three together keeps trading, because no team is looking at the merchant as the unit.

What each team sees, and its blind spot

Team	What it sees	Its blind spot
Merchant risk	The KYB result and the rising chargebacks.	The card-testing pattern and the bot infrastructure behind the spike.
AML / monitoring	The laundering-shaped settlement.	That the merchant identity was stolen and the volume is fake.
Security	The API abuse and the bot traffic.	That it is tied to one funded merchant laundering proceeds.

The joined-up workflow

A sequence to review a high-risk merchant as a single unit rather than three separate alerts.

	Owner	Action	Passes to the others
01	**AML / monitoring**	When settlement or velocity flags a merchant, pull its onboarding record and chargeback trend into the same case.	Merchant ID and laundering indicators to Risk and Security.
02	**Merchant risk**	Re-check the merchant's KYB identity for synthetic or stolen markers and recent identity reuse across the platform.	Identity verdict and linked merchants to AML and Security.
03	**Security**	Match the merchant's API and payment traffic against bot and card-testing clusters.	Linked IPs, bot signature and card-testing hits to AML and Risk.
04	**All three**	Decision point. If identity, laundering and bot signals line up, offboard the merchant, hold settlement, file one SAR, and apply the linked identifiers across the platform to catch sibling merchants.	One merchant case and a platform-wide block.
05	**All three**	Feed confirmed markers back into onboarding and monitoring.	Updated KYB and detection rules.

Signals worth sharing

- → **Merchant risk to AML and Security:** KYB identity markers, document reuse, chargeback trend, linked merchant accounts.
- → **Security to Risk and AML:** bot cluster ID, card-testing fingerprint, API-abuse pattern, source IPs and ASNs.
- → **AML to Risk and Security:** settlement and laundering pattern, velocity anomaly, the cash-out path.

How you'll know it's working

01 Time from the first laundering or card-testing flag to offboarding a bad merchant.

02 Share of bad merchants caught before they reach material volume.

03 Sibling merchants found from one confirmed case.

04 Card-testing traffic blocked at the API rather than settled and charged back.

Where to start

01 Make the merchant, not the transaction, the unit of review for high-risk cases.

02 Run a joint review of card-testing and bot traffic across risk and security.

03 Connect onboarding identity data to live settlement monitoring.

04 Require explainable risk scores you can defend to a scheme or a regulator.

SECTION 04

Remittance firms

The threat picture

Remittance is cross-border, high-volume and low-value per transaction, which is exactly the profile scammers exploit. Romance and investment scams push customer money along international corridors. Mule networks move it onward. Sanctions screening carries real consequences when payments cross borders. And onboarding identity checks are under pressure from deepfakes, often with a compliance team far smaller than the transaction volume would suggest.

Anatomy of an attack: a grooming arc across three corridors

A customer has been in an online relationship for four months with someone he has never met. The story arrives on schedule: a medical emergency abroad, then an investment that can't wait. Over six weeks he sends a series of remittances to three different countries, each one individually unremarkable for a customer who sends money abroad.

The receiving names are mules recruited in the destination markets. The funds are consolidated and moved on, some into crypto, some cashed out through agents. One of the receiving names is a near-match to a sanctions-list entry that the screening system scored just below threshold.

Fraud, if it sees anything, sees a customer sending money abroad with his stated reason on file. AML sees three corridor payments and a near-miss sanctions hit, each handled in isolation. The agent network sees nothing unusual. The pattern, a grooming arc feeding a mule network across three corridors, only exists when the payments are read together.

What each team sees, and its blind spot

Team	What it sees	Its blind spot
Fraud	The customer and the stated reason for sending.	That he is being groomed and the beneficiaries are mules.
AML	The corridor payments and the sanctions near-miss.	That the inbound story is a scam, not a genuine remittance.
Operations / agents	A routine paying-out customer.	The consolidation and onward movement after pay-out.

The joined-up workflow

A sequence to read a customer's corridor activity as one story rather than payment by payment.

	Owner	Action	Passes to the others
01	**Fraud**	On a scam indicator (a sudden new beneficiary abroad, a change in sending pattern, a vulnerable-customer flag), pull the customer's recent cross-border activity into one view.	Sending pattern and beneficiaries to AML.
02	**AML**	Check the beneficiaries against mule indicators and sanctions near-misses across all corridors, not one payment at a time.	Mule and sanctions verdict and corridor links to Fraud.
03	**Fraud**	Where grooming is likely, apply calibrated friction: a hold, a name-check, or a human conversation, sized by the scam-loss risk.	Friction decision and contact outcome to AML.
04	**All three**	Decision point. If corridor activity, mule links and sanctions proximity line up, treat it as a network, freeze onward beneficiaries where possible, and file one SAR covering the arc.	One joined case and beneficiary blocks.
05	**All three**	Feed confirmed mule names, corridors and patterns back into screening and scam detection.	Updated corridor and screening rules.

Signals worth sharing

- → **Fraud to AML:** new-beneficiary events, sending-pattern change, vulnerable-customer flags, the customer's stated reason.
- → **AML to Fraud:** mule confirmation, sanctions proximity, corridor risk, onward movement.
- → **Operations to both:** pay-out anomalies and consolidation behaviour at the receiving end.

How you'll know it's working

01 Share of scam-driven sends stopped before completion rather than reimbursed after.

02 Time from the first scam indicator to reviewing the customer's whole corridor activity.

03 Sanctions near-misses re-scored with fraud context rather than closed in isolation.

04 Mule beneficiaries reused across customers, caught from one confirmed case.

Where to start

01 Pool fraud and AML intelligence on your highest-risk corridors.

02 Calibrate outbound friction using scam-loss data, not a blanket rule.

03 Confirm onboarding withstands deepfake and injected-image attacks.

04 Read a customer's corridor activity together, not payment by payment.

SECTION 05

Crypto and digital-asset firms

The threat picture

Crypto firms operate where transactions are fast and irreversible, which raises the cost of every missed signal. Account takeover and wallet draining run alongside investment scams that end in crypto, on-chain laundering, and sanctions exposure through mixers and flagged addresses. The travel rule adds its own duties. Deepfakes test onboarding here as everywhere, but irreversibility means there is no clawback once the asset has moved.

Anatomy of an attack: a grooming arc across three corridors

A user's account is taken over after his credentials surface in an infostealer dump and the attacker defeats the one-time passcode with a real-time phishing proxy. The attacker changes the withdrawal address and moves the balance out in a single transaction.

On-chain, the funds hop through a mixer and into a cluster of addresses the analytics team has tagged before, some associated with a sanctioned entity. Separately, the same credentials were used days earlier in a low-value test transaction that didn't trip any limit.

Fraud sees an account takeover and a customer complaint. Security sees the credential exposure and the phishing proxy. On-chain analytics sees funds entering a known laundering cluster and brushing a sanctioned address. Three views of one theft, and because off-chain and on-chain data sit in different tools, the withdrawal address is never connected to the takeover in time to act. With irreversibility, time is the whole game.

What each team sees, and its blind spot

Team	What it sees	Its blind spot
Fraud	The account takeover and the customer complaint.	Where the funds went on-chain, and the sanctions exposure.
Security	The credential exposure and the phishing proxy.	That it led to a funded withdrawal into a laundering cluster.
On-chain analytics	The mixer hop and the flagged cluster.	That the source was a specific customer's hijacked account.

The joined-up workflow

A sequence to connect an off-chain takeover to the on-chain trail while there is still time to act.

	Owner	Action	Passes to the others
01	**Fraud**	On an account-takeover or suspicious-withdrawal signal, capture the new withdrawal address and recent login and device changes.	Withdrawal address and account changes to On-chain and Security.
02	**On-chain analytics**	Screen the withdrawal address against known clusters, mixers and sanctioned entities in real time.	Address risk and cluster links to Fraud and Security.
03	**Security**	Check the login for infostealer credentials, real-time phishing-proxy markers and device anomalies.	Credential and device verdict to Fraud and On-chain.
04	**All three**	Decision point. If the signals align, hold the withdrawal where still possible, freeze the account, screen for travel-rule and sanctions obligations, and file accordingly.	One joined case and a held withdrawal.
05	**All three**	Feed the confirmed address, cluster and device markers back into pre-withdrawal screening so the next attempt is blocked before it settles.	Updated screening and withdrawal rules.

Signals worth sharing

- → **Fraud to On-chain and Security:** withdrawal address, account-change events, customer-reported takeover.
- → **Security to Fraud and On-chain:** infostealer credential hits, phishing-proxy markers, device anomalies.
- → **On-chain to Fraud and Security:** cluster and mixer links, sanctioned-address proximity, travel-rule counterparty data.

How you'll know it's working

01 On-chain to Fraud and Security: cluster and mixer links, sanctioned-address proximity, travel-rule counterparty data.

02 Time from a takeover signal to holding or freezing, measured in minutes.

03 Sanctioned-address proximity caught pre-withdrawal rather than reported after.

04 Repeat addresses and device markers blocked at the next attempt.

Where to start

01 Connect on-chain and off-chain risk into one real-time view.

02 Set phishing-resistant MFA and deepfake-resistant onboarding as baseline.

03 Screen withdrawal addresses before execution, given irreversibility.

04 Require explainability on AML and sanctions decisions.

SECTION 06

From shared intelligence to shared action

The threats in this guide look different by sector, but the pattern is remarkably consistent. Fraud, AML and Cyber each hold part of the picture; the opportunity is to connect those signals early enough to change the outcome.

Much of that starts with people and process: clearer ownership, better hand-offs and agreement on which intelligence should move between teams. Technology then determines how easily those ways of working can happen at scale.

A shared view is only valuable if teams can explore the data behind it, identify meaningful connections, turn what they learn into better detection, and explain the decisions that follow.

Where Fortify fits

Fortify is an AI-native AFC platform proven to help teams move from fragmented data and detection to faster, more adaptive financial crime operations. Unlike legacy tools that only flag risk, Fortify enables users to act quickly.

It lives within your existing data warehouse, so your data stays where it belongs. From there, teams can:

- → **Explore data using natural language,** giving users a simpler way to interrogate the information available to them.
- → **Build and test detection rules or retrain models in minutes,** without waiting in an engineering queue.
- → **Turn new intelligence into better detection faster,** so lessons from investigations and emerging threats can feed back into your controls.
- → **Maintain a complete audit trail at the click of a button,** so decisions can be understood, evidenced and reviewed.

Put it to work on your own data

The examples in this guide are illustrative. Your threats, data and controls will be specific to your organisation.

Fortify can help you envision what a more connected, adaptive approach could look like using the data and financial crime use cases you already have, while working within your existing technology environment.

Speak to our experts →