

Guide:


Can You See the Network Behind Your Scam Losses?

Examining the five layers to effectively detect, link, and block scams




INTRODUCTION

Introduction



Your scam loss figure is accurate, but it describes the last step of a longer chain. Scam losses are counted at the payment, but they are stopped underneath the payment, in layers that a reported case never shows. A reported case is one row in your data: a clean payment recorded with the correct credentials and device, and one beneficiary account that you identify days or months after the money moved. The **Scam Iceberg** pictures what that row leaves out. Beneath the reported payment sit five layers (Figure 1).



- 1 Mule account intake.** Criminals open or buy accounts, known as **mule accounts**, to receive the money.
- 2 Identity assembly.** They build identities that pass verification.
- 3 Device and session reuse.** They run many accounts from the same devices and sessions.
- 4 Money movement and cash-out.** They move the funds onward to a cash-out point.
- 5 Network coordination.** A criminal network coordinates the whole chain across many institutions.

Figure 1. The Scam Iceberg: five layers beneath one reported scam payment.

Scam losses are counted at the payment, but they are stopped underneath the payment, in layers that a reported case never shows.

For scams in which the customer authorizes the payment, liability has increasingly shifted to the sending institution, so the loss lands on your books. [UK Finance's Annual Fraud Report 2026](#) shows that UK banks reimbursed £354.3 million to victims of those scams in 2025, about 61% of their total losses.

A reimbursed case then closes as a write-off while the criminal network behind it keeps its accounts, identities and devices and can reuse them against you. Every new case reaches your team as one row with no link to the others in the queue.

Seeing how your cases connect is how you find the mule accounts, networks and entry points that need stronger controls before the next loss.



Loss value alone cannot show how your scam cases connect, and it measures only what reporting captured.

For 2025, the Consumer Federation of America's [The Scam Economy](#) compared \$20.8 billion in US consumer fraud losses reported to the Federal Trade Commission (FTC) with an implied \$148.2 billion in actual losses.

Your own case records have the same blind spot: a scam that was never reported never becomes a case. Link analysis shows how the cases you did confirm connect to each other, and those connections can point to activity that never reached your caseload.

US consumer fraud losses, 2025

Reported to the FTC

\$20.8 billion



Implied actual losses

\$148.2 billion



Consumer Federation of America, The Scam Economy

Are you identifying where your scam losses come from, and stopping them there? Let's walk through the five layers of the Scam Iceberg to examine how your institution detects, links, and stops scams at each one.

EXISTING CONTROLS

What your existing controls already cover

You can only address gaps once you know what is already covered, so credit what your current controls already do before you look for gaps. Your **gap list** then holds only the layers your controls do not yet cover.

Payee verification, a control you may already run, shows the method: for each control, state what it covers and where that coverage stops, because a gap begins where coverage ends:



What payee verification covers

[Pay.UK](#) reports that over 300 organizations have implemented **Confirmation of Payee**, an account name-checking service for UK domestic payments, and that more than 2 million checks are completed every day.

Pay.UK describes the service as designed to help reduce misdirected payments.

The European Union now requires payment service providers to offer payers a **verification of payee service** before a credit transfer is authorized.

300+

organizations have implemented Confirmation of Payee

2 million+

checks completed every day


Pay.UK




Where payee verification stops

The check compares the name a payer entered with the account holder's name. It cannot tell you whether the payer was given that name for an honest reason, so a mule account opened in the name the victim was told to expect returns a clean match and the payment goes through.

How payee verification works, and where it stops



Record three things about your own controls before you go further, so that your gap list starts from what is already covered.



The third answer decides how much of this work you can finish with a **join**, a query that links records from different systems, and how much needs new work.

Because payee verification confirms only that a name matches an account, how the receiving account behaves is the first gap to examine, and Layer 1 starts there.

LAYER 1 OF 5

Layer 1: mule account intake

The receiving account is the most testable layer in the chain, because every scam payment lands in one. [UK Finance's Annual Fraud Report 2026](#) recorded £576.4 million in **authorized push payment (APP)** scam losses across 248,070 cases in the UK in 2025.

In an APP scam the customer authorizes the payment, so every one of those cases ended with money arriving in a receiving account at some institution, and that institution can test how the account behaves.



Figure 1, repeated. The Scam Iceberg with Layer 1, mule account intake, highlighted.




What hides here

Criminals open or buy accounts to receive scam payments. Some of those accounts stay dormant before their first large credit, and some are opened under an identity assembled for the purpose.



What surfaces it

Monitoring on the inbound side of an account shows what outbound monitoring cannot. Useful signals include how fast credits arrive compared with the purpose given at account opening, how quickly money leaves after a credit arrives, and the ratio of inbound to outbound activity over the account's life. Judge dormancy against that stated purpose, because an account that sat quiet for months can pass every age-based test and differ only in its first significant credit.





£576.4 million

in APP scam losses



248,070

cases, in the UK in 2025

UK Finance, Annual Fraud Report 2026

Three checks to run

- ☐ Take your confirmed mule accounts from the last year. What share showed unusually fast inbound credits before you confirmed them, and how many days earlier was the signal available?
- ☐ For what percentage of your accounts is a purpose recorded at opening that your monitoring can compare against observed flow?
- ☐ When a beneficiary account is confirmed as a mule account, does anything automatically check whether other accounts share its device, address or phone number?

What a gap looks like

You learn that an account was a mule account only from an inbound fraud report sent by another institution, with no internal signal before the report.

Who to involve

Bring in whoever owns account monitoring and whoever owns onboarding. Inbound monitoring protects other institutions' customers as well as your own, so the business case rests on repeat exposure and regulatory duty as well as direct loss avoidance.

LAYER 2 OF 5

Layer 2: identity assembly

Every receiving account starts with onboarding, where it has to pass identity verification, so onboarding is the checkpoint criminals have to get past before they can take scam payments.

The FTC's [Consumer Sentinel Network Data Book 2024](#) shows that US consumers reported losing \$2.09 billion to fraud by bank transfer and payment in 2024, the largest single payment method reported, and every payment of that kind lands in a receiving account.




Figure 1, repeated. The Scam Iceberg with Layer 2, identity assembly, highlighted.



What hides here

Criminals build **synthetic identities**, which are identities made from a mix of real and invented elements, to pass verification without impersonating a named victim. They can leave an identity unused while it builds a history, since each small credit line repaid on schedule makes the next application look better founded.



What surfaces it

Comparing applications with each other finds what checking each one alone misses. Each element of a synthetic identity, such as a name, an address or a phone number, can pass verification on its own. What is missing is a consistent relationship between those elements over time, and only a comparison across applications can reveal that missing relationship.

↔ **\$2.09 billion**

reported lost by US consumers to fraud by bank transfer and payment in 2024, the largest single payment method reported

FTC, Consumer Sentinel Network Data Book 2024

Three checks to run

- ☐ Can you query how many approved applications in the last two years share an identity element with another application while differing on the rest?

- ☐ Are losses from accounts that never disputed anything booked as fraud losses or as credit losses? Collections holds the answer.

- ☐ Does your onboarding decision record keep the data on each identity element, such as name, address and phone number, that you need to answer the first check, or only the final decision?


What a gap looks like

Onboarding and fraud monitoring share no common identifier, so the link between account opening and later payment fraud cannot be measured.

A synthetic identity written off as a credit loss carries no fraud mark, so it can be used again, and the accounts linked to it are never traced.

Who to involve

Ask collections directly. Where synthetic identity losses are booked as credit losses, the fraud function never sees them, the credit function has no reason to call them fraud, and the total sits in a category nobody measures against a detection roadmap.



LAYER 3 OF 5

Layer 3: device and session reuse

A device link helps only if it is available on the same timescale as the scam, and scams move quickly. The [Global Anti-Scam Alliance \(GASA\)](#) reported in its 2024 *Global State of Scams* report that nearly half of scams are concluded within 24 hours, from first contact to payment.

Device links found only during case review arrive after the payment has completed and explain a single case, while a link your team can see in real time connects that case to other cases that share the device.




Figure 1, repeated. The Scam Iceberg with Layer 3, device and session reuse, highlighted.

24 hours



Nearly half of scams are concluded within 24 hours, from first contact to payment

GASA, 2024 Global State of Scams



What hides here


A single criminal network can run many accounts from shared infrastructure, such as the same devices and sessions.




What surfaces it

Device and session linkage connects accounts that show no other relationship, for example two beneficiary accounts opened from the same device.

DEVICE DATA AND PAYMENT MONITORING, NOT JOINED



WITH A JOIN



Three checks to run

- ☐ For your confirmed scam cases, how many distinct accounts shared a device identifier with the beneficiary account, and was that visible at the time or only in hindsight?

- ☐ Which team holds your device data, and can payment monitoring query it?

- ☐ How long after a device is linked to a confirmed scam case is that device checked against your open accounts?

What a gap looks like

Device intelligence is held for login and authentication decisions but is never joined to beneficiary accounts or payment monitoring.

Who to involve

If device data sits with an authentication or digital team and payment monitoring cannot query it, nothing needs to be bought. The join needs to be built, and someone needs to own the request. Name both owners next to the gap on your gap list.

LAYER 4 OF 5

Layer 4: money movement and cash-out

[ACI Worldwide's Scamscope report](#) projects that real-time payments will account for 80% of APP scam value across six leading real-time payment markets (the US, UK, India, Brazil, Australia and the UAE) by 2028.

ACI adds that scammers exploit the immediacy of these payments to steal funds before they can be traced. If ACI's projection holds, most scam value will move on rails built for speed, so what happens to the money after it arrives, the part a reported payment does not show, matters more.

Once scam money reaches a beneficiary account, it can keep moving through further accounts, institutions and rails. The path the money takes forms a trail, and the reported payment is only the first step along it. The loss is counted at that first step, but the scam is understood in full only by following the money beyond the first account.

Without that view, your team reimburses each loss as it comes in and never addresses the fundamental problem, the scam network and its accounts, which keep producing new losses.




Figure 1, repeated. The Scam Iceberg with Layer 4, money movement and cash-out, highlighted.



80%

of APP scam value is projected to move on real-time payments by 2028, across six markets

ACI Worldwide, Scamscope




US, UK, India, Brazil, Australia, UAE

Following that trail is what Layer 4 of the Scam Iceberg covers: money movement and cash-out. To examine this layer, test whether your records can follow the money beyond the first account, and how far they get toward a cash-out point.

Every step you can follow brings more of the scam network into view: the accounts that carried the money, and the point where the network collects it. Each of those is a place where your controls can be strengthened against the network itself.

The same trail matters beyond fraud losses, because your financial crime controls need to see where the money from a scam goes, and a trail you can follow is a trail you can explain to an examiner.

Following the scam trail





What hides here

Criminals split funds across accounts, institutions and rails on the way to a cash-out point.



What surfaces it

Fund-flow tracing follows the value across every account it passes through.

Three checks to run

- ☐ What is the median time between funds arriving in a beneficiary account and leaving it, and does any control act inside that window?

- ☐ Can you trace a scam payment more than one **hop** (one transfer from one account to the next), and how many analyst hours does one trace take?


- ☐ For your confirmed cases, how often can you see where the money went after the beneficiary account, and how do you find out?

What a gap looks like

Tracing stops at the first hop, or each trace takes analyst hours that your caseload cannot spare.

Who to involve

If fund-flow tracing sits with financial crime investigations, ask that team what a one-hop trace and a multi-hop trace cost today, and use the answers to size the gap for your committee.



LAYER 5 OF 5


Layer 5: network coordination

Your view stops at your own perimeter and the criminal network behind a scam does not, which makes this the one layer no institution can fully see on its own. In [United States of Scams: The Financial and Emotional Fallout](#), Gallup and the Stop Scams Alliance estimate that US adults lost \$68 billion to scams in 2025, with about 15 million adults victimized, and that 56% of reported scams cost \$500 or less.

Amounts that small can sit under the thresholds a rule is written around, and a rule that fires on one payment cannot show that many such payments came from the same network.



Figure 1, repeated. The Scam Iceberg with Layer 5, network coordination, highlighted.



\$68 billion

estimated US scam losses in 2025

Gallup and Stop Scams Alliance, United States of Scams

about 15 million adults victimized

56% of reported scams cost \$500 or less



What hides here

A criminal network can work against many institutions at once and reuse accounts, identities and devices that each institution closes locally and nobody closes globally.



What surfaces it

In a **consortium**, participating organizations contribute anonymized fraud signals to a shared pool, so a bad actor flagged by one participant can help protect the others, and the pool can surface connections such as shared devices, addresses or payment methods that one institution's own transaction history would not reveal.

One example is FraudNet's Global Anti-Fraud Network, which pools anonymized fraud and risk signals contributed across participating organizations, so any one customer benefits from patterns the whole network has observed, beyond its own history.

Three checks to run

- ☐ When you close a mule account, is anything about it shared outside your institution, and through what mechanism?

- ☐ How often does a newly confirmed mule account turn out to have been reported by a peer institution earlier?

- ☐ Do you receive any external signal on a beneficiary account before the first credit, or only after a fraud report?

What a gap looks like

The gap appears as repeat exposure: the same criminal network returns with accounts you never had visibility to close.

Who to involve

Ask whoever represents your institution in industry information-sharing arrangements what you share today and what you receive in return.



Turn your gap list into a plan

By the end of these checks you know where you may have gaps. Take the findings to your risk committee to start the conversation about the specific gaps that should most immediately be addressed, and when.

You may find, for example, that losses from synthetic identities are written off as credit losses. The fraud team never counts them, so the work of catching them can go under-funded, and the same identity can be used again. Or, you may find that tracing stops at the first hop, so you cannot follow a money trail.

Sort the gaps by the kind of work each one needs. Where data was never collected, the gap is a roadmap item. Where data is collected but sits in separate systems, the gap is work you can schedule this quarter. Sorting gaps this way separates what you need to buy from what you already own and have not connected.

Join the data you hold, and add signals from beyond your own history. **FraudNet's Customer Risk Intelligence solutions** bring identity, device and behavior together in one risk score. Its Global Anti-Fraud Network shares collaborative insights across industries, so a mule account flagged elsewhere can be known here.

Four capabilities support the analysis and the work that follows:



Unified, Entity-First View of Risk

Identity, device and behavior come together in one risk score.



Real-Time Transaction Monitoring

Payments via digital, card, ACH, wire & branch channels are screened as they happen.



Global Anti-Fraud Network

Collaborative insights are shared across industries.



Behavioral Scam Detection

Behavioral detection identifies scam typologies to block fraud in real-time, not after the fact.

An institution with visibility across its whole chain can better identify scam networks and act immediately.

To see how FraudNet can help your team address the gaps and pain points you find, **book a meeting**.

[Learn more](#)