



REPORT

Identity Fraud Report 2026

How coordinated fraud networks
target remote onboarding



Table Of Contents

01	03
----	----

Executive Summary

03	06
----	----

Attack Typologies Against Remote Identity Verification

05	21
----	----

Identity Fraud Across Industries

07	36
----	----

Continuous Identity Verification

09	38
----	----

The Identity Verification Solution to Prevent Fraud Risk

02	04
----	----

Introduction

04	13
----	----

The Rise of Coordinated Fraud Networks

06	34
----	----

Building Identity Resilience

08	37
----	----

Conclusion

10	40
----	----

References

Executive Summary

Generative AI has lowered the technical barrier and cost of producing convincing synthetic content

Attackers can now create and modify identity artefacts at greater scale. Fraud attempts increasingly combine multiple attack methods, including manipulated documents, synthetic identities, and biometric spoofing techniques, and Shufti's data shows that some fraudulent attempts share reused verification data across multiple verification events.

WHAT IT COSTS WHEN THESE ATTEMPTS GO UNIDENTIFIED

FOREX AND IGAMING

Losses to multi-trading accounts, multi-player accounts, and jurisdiction hopping.

CRYPTO EXCHANGES AND FINTECHS

Non-compliance risk where fraud rings and mule networks go undetected.

LAW ENFORCEMENT

These patterns create additional challenges in identifying and disrupting cross-border fraud activity.

22.49%

Identity fraud reached **22.49 percent of all verification requests in crypto and exchange**, the highest of any sector measured. The findings highlight how sectors with high volumes of remote onboarding remain attractive targets for identity fraud attempts.

70

The largest linked cluster observed connected **70 identities across 13 devices**, based on shared fraudulent documents, devices, and IP-related signals.

38 sec

In the fastest cross-border network, the same operation signed up from a different jurisdiction just **38 seconds later**, a timing pattern that is inconsistent with normal physical movement between jurisdictions.

Networks,
not actors

Deepfakes are not only exploited by fraudsters to bypass a KYC check. These typologies also appear in **cross-border activity where one forged document is reused across many identities**.

Introduction

Every remote interaction depends on one assumption

That the person on the other side is who they claim to be. Opening an account, drawing down a loan, funding a trading account, or registering to play, all rest on it.

This assumption is increasingly being targeted as remote identity verification systems become a focus for criminals, fraud networks, organised crime groups, and even state-linked threat actors.

These groups are using AI-powered tools, synthetic identities, and deepfake techniques to impersonate legitimate individuals and manipulate identity checks. The FBI has documented North Korean IT worker campaigns in which operatives used stolen identities, fraudulent accounts, and AI-generated video to conceal their true identities and gain trusted access to organisations.¹

These cases demonstrate how identity verification systems can be exploited when attackers manipulate the signals used to establish trust.

An AML obligation and a loss line at the same time

For the businesses doing the remote onboarding, fraud is an AML obligation and a loss line at the same time. Regulated firms must know who they are onboarding and be able to evidence it, while the same fraudulent accounts convert directly into money out the door.

Where the exposure is sharpest

The exposure is sharpest in iGaming, forex, and lending, where one operator running accounts under fabricated or stolen identities can farm sign-up bonuses at scale, hold opposing positions across multiple trading accounts, or draw credit against an identity that does not exist.

Losses in these sectors run into millions of dollars, and they surface as chargebacks, unrecoverable trading losses, written-off credit, and licensing consequence long after each account cleared onboarding individually.

Methodology

This report is based on Shufti's internal identity verification data for the first half of 2026, covering eleven industries: crypto, fintech, banking, gaming and iGaming, payments, lending and investment, telecom, e-commerce and marketplaces, forex, insurance, and healthcare. Figures are derived from requests declined after a associated fraud was confirmed. The report sets out the relationships observed and their cluster sizes, and attributes no attempt to a named individual, organisation, or jurisdiction.

The first analysis


Tracks the AI-powered fraud trend: the share of declined attempts using an AI-generated or AI-manipulated document or biometric, split between deepfake documents, synthetic identities, injected video, and face swaps.

The second analysis

Counts only verification events where the identity document was found to be fraudulent, manipulated, or a deepfake, then links those events through shared attributes. It shows not only that manipulated artefacts are more common, but that the same artefacts are reused across devices, IP addresses, and infrastructure to bypass KYC.

THREE PATTERNS ACCOUNT FOR MOST OF WHAT SURFACES

- A manipulated document already declined for failing the originality check is resubmitted through a different device or IP address.
- A single device or shared IP address is used to submit several different identity documents.
- Shared devices, IP addresses, and identity attributes recur across a group of attempts, including in the cross-border cases and resemble the pattern of historic fraud.



Attack Typologies Against Remote Identity Verification

Presentation Attacks

A presentation attack, as per ISO/IEC 30107-1, is a presentation to the biometric capture subsystem with the goal of interfering with the operation of the biometric system, countered by presentation attack detection. The manipulated item, a deepfake, an altered document, or a mask, is the presentation attack instrument.

The attack is the act of presenting it to the capture device, and in remote identity verification that applies across the capture channel to both the face and the document. These map to the detection categories quantified per industry in Section 5.

Categorised by what is presented and how

- **AI-generated and deepfake presentations:** a synthetic or face-swapped face, or an AI-fabricated document image.
- **Altered-document presentations:** a genuine document modified by photo substitution, data change, or physical tampering.
- **Replay and screenshot presentations:** a previously captured image or video re-presented to the camera.
- **Liveness and physical-material presentations:** masks, cut-outs, or other artefacts to defeat liveness detection.

Presentation Attack Instruments and How Each Is Presented

Every presentation attack, whatever the artefact, has to travel the same physical route to reach a decision. The instrument is prepared, it is held up to the lens, the camera captures it as if it were a live subject, and presentation attack detection is the single control standing between that artefact and an approved account.

Because all four instrument types converge on one channel, the camera, detection at that point covers the whole class. That is also why an attack that never reaches the lens, covered next, defeats these controls entirely.




Figure 1. The Presentation Attack Detection

Injection Attacks

Nothing is ever presented to a real sensor

Injection attacks skip the physical presentation stage entirely and insert manipulated images, biometric data, or video streams straight into the verification pipeline, using virtual cameras, emulators, camera-feed replacement, API-payload manipulation, and frame injection.

Because they never present to a real sensor, they defeat controls built only for presentation attacks and stay harder to catch with conventional liveness checks. Injection therefore becomes a software-integrity problem as much as a document one, and the strongest tell is often the infrastructure the attack rides on rather than the artefact itself.

The principal methods include deepfake video injection and injected selfie images, virtual-camera attacks and camera-feed replacement, and API-payload manipulation and frame injection. The question shifts from what the image looks like to whether the capture channel itself can be trusted.




Figure 2. The Injection Attack Detection

Infrastructure-Based Identity Fraud

Infrastructure-based identity fraud refers to coordinated attacks in which fraudsters repeatedly leverage shared technical infrastructure, including device emulators, spoofed devices, VPNs, proxy networks, controlled IP addresses, and reused identity assets, to run identity verification attacks at scale.

Multiple fraudulent identities sharing common infrastructure can be a strong indicator of coordinated activity, although additional signals are required to determine the nature of that relationship.

Deepfakes, synthetic faces, forged documents, and injected video work as interchangeable inputs, so a forged document in one attempt can be paired with a face swap in the next and a synthetic profile in a third, all run through the same shared devices, IP addresses, and reused assets.

The data shows patterns in how these instruments appear together. A cluster of identities sharing common devices, addresses, or identity assets can indicate coordinated activity, while determining the specific actors behind that activity requires further investigation.

UNITED KINGDOM

FCA on Combating the Mule Networks

In its multi-firm review of the proceeds of fraud, published on 19 October 2023, the UK Financial Conduct Authority observed firms onboarding multiple customers who shared the same device without a clear explanation, which it noted can characterise a money mule herder controlling accounts opened in different names. It suggested firms consider device profiling, geolocation, and behavioural biometrics within onboarding and monitoring, while allowing for legitimate explanations such as members of the same household.

HONG KONG

The HKMA on Network Analytics

In a report published in May 2023, the Hong Kong Monetary Authority set out how analysing accounts as a connected network, rather than one at a time, reveals mule-account relationships that single-account review does not. It reported that around 60 percent of Hong Kong retail banks were deploying network analytics, more than double the level three years earlier.

The background is a blurred office scene. In the foreground, a person's hands are visible typing on a keyboard. Several large computer monitors are on the desk, displaying various data and charts. In the background, other people are working at their desks, and the office is filled with typical equipment like coffee cups and papers.

Fraud Trends Across Various Identity Verification Artefacts

Identity Verification Artefacts

Four instruments do most of the work across the estate

These instruments are the documents, faces, and identities presented to the verification step, distinct from the infrastructure that carries them. Shufti's Deepfake Identity Fraud Index Report 2026 sizes the AI-driven share of that activity and splits it across four instruments.⁴

Deepfake Document Fraud

A genuine document altered by photo substitution, changed data fields, or a reprinted template, through to a wholly counterfeit identity page. The same forged document is reusable across many applications, which is why it is the artefact most likely to reappear inside a coordinated network.

80.10%

of AI-enabled fraud in the first half of 2026, the dominant attack instrument by a wide margin.

Synthetic Identities

Real and fabricated details combined, a valid identifier attached to an invented name and history, into a person who does not exist but looks legitimate on paper. Because a synthetic identity is assembled rather than stolen, it clears checks that only test whether individual data points are internally consistent.

12.31%

of AI-enabled fraud, built to pass onboarding and then age quietly.

Injected Videos

A pre-recorded or generated stream fed straight into the verification pipeline through a virtual camera or manipulated feed, bypassing the camera altogether. Nothing is presented to a real sensor, so the question shifts from what the image looks like to whether the capture channel itself can be trusted.

4.01%

of AI-enabled fraud, fed into the pipeline without reaching a real sensor.

Face Swaps

The applicant's face replaced with someone else's, or with a synthetic composite, at the biometric and liveness stage. The tell sits in the mismatch between the presented face, the document portrait, and the way a real face behaves under liveness prompts, rather than in any single frame.

3.58%

of AI-enabled fraud, presented at the biometric and liveness stage.



The Rise of Coordinated Fraud Networks

Coordinated Fraud Networks

The instrument-level fraud is only the entry mechanism. The exposure sits in the repetition.

Organised fraud is increasingly enabled by interconnected criminal ecosystems where technology, specialised services, and supporting infrastructure allow illicit activities to operate at greater scale. UNODC's Transnational Organized Crime Threat Assessment for Southeast Asia 2026 describes how organised crime groups are adopting technology-enabled models and relying on interconnected services and capabilities to expand their operations.⁵ This shift shows how fraud operations can move beyond isolated incidents towards structured activities supported by wider digital ecosystems. Shufti's verification data reflects this broader trend at the identity layer, where fraudulent verification attempts can share identity artefacts, devices, and infrastructure signals, creating connections between seemingly separate attempts.

A forged or AI-generated document, a face swap, an injected video, and a synthetic identity are interchangeable inputs, and both individual fraudsters and organised networks combine them to defeat KYC repeatedly rather than once.

An individual may cycle through several instruments to open a run of accounts that a single-customer control was meant to prevent, farming sign-up bonuses or evading self-exclusion across multi-player accounts in iGaming, or holding multi-trading accounts in forex to abuse promotional terms and spread activity beneath position limits.

A coordinated network does the same at scale, distributing many identities across accounts, jurisdictions, and payment routes so that no individual account carries the pattern.

By connecting identities through reused fraudulent documents, matching faces, shared devices, and controlled IP addresses, identity fraud analytics can reveal relationships between apparently separate verification attempts, turning individual signals into a broader picture of potentially coordinated activity. What supervisors may identify through investigations and case work, Shufti can observe through patterns in live verification traffic. The figures that follow represent observed relationships within production data across a large sample rather than conclusions about specific actors.

Multi-Identity Account Abuse

One person uses several genuine, stolen, synthetic, or manipulated identities to open accounts that appear unrelated, usually to repeat a benefit or bypass a customer-level control such as a bonus, a lending limit, a self-exclusion rule, or an account suspension.

The exposure is sharpest in iGaming, where a single actor can farm sign-up bonuses and slip past self-exclusion across dozens of accounts. The UK Gambling Commission documented one such case, in which more than 5,900 identity and financial-document scans were used to generate nearly £80,000 through fraudulent gambling accounts.⁶

SHUFTI'S NETWORK FRAUD DATA FOR 2026 SHOWS

- **65.68 percent** of cross-attempt matches trace to reuse of the same fraudulent identity document, the strongest single signal that several attempts may be related.
- The largest linked identity cluster observed connected **70 identities across 13 devices**, based on shared document, device, and infrastructure signals, with one device alone appearing across 16 separate verification requests.
- **2.01 percent** of attribute-sharing fraudulent attempts span more than one country, the cross-border tail where identity fraud meets money laundering.

Pattern of Network Fraud Through Shared Attributes

Which shared signal actually indicates a common operator? People legitimately share devices and networks, so each signal carries different weight.

01 Shared fraudulent documents

Reuse of a document already declined as forged, manipulated, or AI-generated. The strongest matching signal, and the dominant thread at 65.68 percent.

02 Shared controlled IP addresses

The same network origin behind attempts that otherwise look unrelated. Corroborating rather than conclusive on its own.

03 Shared devices

One device carrying several identities. When it appears alongside a reused document, the two signals confirm each other and AI-assisted fraud makes such attempts inexpensive to produce at scale.

Composition of shared-attribute matches





Figure 3. Composition of shared-attribute matches between fraudulent attempts.

The Largest Observed Identity Fraud Network



70

Identities linked in one cluster

13

DEVICES

16

EVENTS, ONE
DEVICE

5

EVENTS,
ONE IP


The largest connected cluster observed in the first half of 2026, linked through shared fake documents, devices, and IP-related indicators.

Figure 4. One connected identity cluster. Each node on the ring is a device, each dot a verification event carried on it. The distribution beyond the sixteen-event hub is illustrative.

Share of Different Fraud Clusters

The overwhelming majority of linked clusters involve just two identities, consistent with related account activity rather than unrelated users coincidentally sharing attributes.

Where those identities sit on different devices, the pattern is characteristic of multi-accounting, particularly relevant to iGaming and to forex and online trading.



Each ring is one linked cluster. Dots are identities, the centre marks the shared attribute connecting them.

Cross-Border Shared-Attribute Fraud

Shufti's identity fraud data from the first half of 2026 shows that 2.01 percent of fraudulent verification attempts that share identity or infrastructure attributes span more than one country, a strong indicator of cross-border coordinated activity.

The pattern represents shared identity or infrastructure signals appearing across jurisdictions. The hardest networks to disrupt are those that do not stay in one country. When a chain crosses borders, no single institution or regulator sees the whole of it, and an identity-verification gap becomes a money-laundering one.

The timing reduces the likelihood that these attempts represent independent activity or ordinary customer movement. Within these cross-border groups the typical gap between activity in one country and the next is 9 minutes 33 seconds, and the shortest groups run their full course in well under a minute. Activity that surfaces in two countries minutes apart may indicate shared infrastructure or coordinated activity rather than normal physical movement between jurisdictions.

This kind of coordination follows the pattern of mule networks operating across borders, account holders recruited or fabricated in one jurisdiction, run through shared infrastructure, and used to move value into another. It is also the mechanism behind cross-border recruitment fraud, where a fabricated candidate identity is assembled in one country, presented to an employer in another, and used to obtain payroll access, credentials, and systems entry without the candidate ever being physically met.

A single cross-border cluster can carry significant laundering exposure and is where these cases intersect with AML and sanctions obligations, so identifying it at onboarding materially changes the outcome.

Cross-Border Timing and Concentration


Time between one country and the next

38 sec
FASTEST OBSERVED

9 min 33 sec
TYPICAL GAP

Activity jumps between countries in under 10 minutes typically, and in as little as 38 seconds — too fast to be physical travel.

Share of activity by jurisdiction



A dark, moody photograph of two individuals wearing hooded sweatshirts in a server room. They are looking at a laptop screen. In the background, there are large server racks with multiple monitors displaying data. The lighting is dim, with a blueish-green glow from the screens.

Identity Fraud Across Industries

Where Identity Fraud Concentrates

Exposure ranges from 22.49 percent of all verification requests in crypto and exchange down to 4.24 percent in banking, tracking the sectors whose remote onboarding volume is highest.




Figure 6. Identity fraud as a share of all verification requests, by industry, first half of 2026.

Fraud in the Crypto Sector

Altered documents lead the sector at 10.10 percent, ahead of failure to prove liveness at 3.84 percent.

INDUSTRY FRAUD EXPOSURE

22.49%

of all verification requests in Crypto Sector were identity fraud in the first half of 2026

Identity fraud in the crypto sector reached 22.49 percent of all verification requests in the first half of 2026. The Financial Action Task Force identifies virtual asset service providers as exposed to identity fraud at onboarding, where applicants submit forged or digitally edited identification documents and doctored photographs to defeat KYC checks. FATF’s red-flag work documents cases in which offenders used falsified identification to circumvent KYC at multiple exchanges and routed large volumes of stolen virtual assets through providers with weak controls.⁷

ATTACK TYPOLOGY BREAKDOWN




Figure 7. Crypto Sector: detection rate by attack instrument.

Shufti’s data mirrors that exposure. Altered documents reach 10.10 percent, the highest single-instrument rate anywhere in this report, followed by failure to prove liveness at 3.84 percent, replay and screenshot attempts at 3.16 percent, and AI and deepfake images at 2.62 percent. Crypto is also one of only two sectors where liveness failures outrank replay attempts, which suggests attackers working the biometric step as hard as the document one.

Fraud in the Fintech Industry

Altered documents lead the sector at 3.79 percent, ahead of failure to prove liveness at 2.47 percent.

INDUSTRY FRAUD EXPOSURE

18.36%

of all verification requests in Fintech were identity fraud in the first half of 2026

Identity fraud in fintech reached 18.36 percent of all verification requests in the first half of 2026. FinCEN has warned that criminals are using generative AI to fabricate identity documents, photographs, and videos that defeat the customer identification controls underpinning remote, digital account opening, the operating model fintechs are built on. Its alert records accounts being opened with suspected

GenAI-produced identities and then used to receive and launder the proceeds of other fraud.⁸ In Shufti’s data the sector is led by altered documents at 3.79 percent, ahead of failure to prove liveness at 2.47 percent, replay and screenshot attempts at 1.61 percent, and AI and deepfake images at 1.41 percent. Fintech is the other sector where liveness failures outrank replay attempts.

ATTACK TYPOLOGY BREAKDOWN



Figure 8. Fintech: detection rate by attack instrument.

Fraud in the Banking Industry

Replay and screenshots lead the sector at 1.23 percent, ahead of altered documents at 0.77 percent.

INDUSTRY FRAUD EXPOSURE

4.24%

of all verification requests in Banking were identity fraud in the first half of 2026

Identity fraud in banking reached 4.24 percent of all verification requests in the first half of 2026. FinCEN’s analysis of Bank Secrecy Act reporting shows identity exploitation is pervasive across banking, driven mainly by impersonation using another person’s or a fabricated identity. Depository institutions filed the largest share of identity-related suspicious activity reports, around 1.3 million filings covering about \$201 billion, which places account-opening and identity-verification

processes at the centre of the sector’s exposure.⁹ In Shufti’s data, replay and screenshot attempts lead banking at 1.23 percent, ahead of altered documents at 0.77 percent, AI and deepfake images at 0.50 percent, and failure to prove liveness at 0.47 percent. Banking is the only sector in the report where altered documents do not lead, which may reflect differences in onboarding controls, customer behaviour, fraud patterns, or detection approaches.

ATTACK TYPOLOGY BREAKDOWN




Figure 9. Banking: detection rate by attack instrument.

Fraud in the Gaming & iGaming Industry

Altered documents lead the sector at 4.43 percent, ahead of ai and deepfake images at 2.50 percent.

INDUSTRY FRAUD EXPOSURE

12.45%

of all verification requests in Gaming & iGaming were identity fraud in the first half of 2026

Identity fraud in gaming and iGaming reached 12.45 percent of all verification requests in the first half of 2026. The UK Gambling Commission’s money-laundering and terrorist-financing risk assessment flags remote gambling for false or stolen identity documentation used to bypass controls, and for customers who are not physically present for identification. It rates mule accounts as high risk and singles

out inadequate KYC, including checks completed only at the point of withdrawal rather than at onboarding.¹⁰ Shufti’s data shows altered documents leading the sector at 4.43 percent, in line with a document-integrity exposure at account opening, followed by AI and deepfake images at 2.50 percent, replay and screenshot attempts at 2.47 percent, and failure to prove liveness at 0.60 percent.

ATTACK TYPOLOGY BREAKDOWN




Figure 10. Gaming & iGaming: detection rate by attack instrument.

Fraud in the Payments Industry

Altered documents lead the sector at 7.90 percent, ahead of replay and screenshots at 3.73 percent.

INDUSTRY FRAUD EXPOSURE

14.87%

of all verification requests in
Payments were identity fraud in the
first half of 2026

ATTACK TYPOLOGY BREAKDOWN




Figure 11. Payments: detection rate by attack instrument.

Identity fraud in payments reached 14.87 percent of all verification requests in the first half of 2026. FATF’s study of cyber-enabled fraud found that the payment service providers most involved were around 200 times riskier than other institutions, specifically because of poor identity verification and transaction monitoring. Criminals set up and control accounts using stolen or falsified identities,

which makes the resulting money-laundering activity harder to trace.¹¹ Shufti’s data reflects that document weakness. Altered documents lead payments at 7.90 percent, the second highest single-instrument rate in the report, ahead of replay and screenshot attempts at 3.73 percent, AI and deepfake images at 1.67 percent, and failure to prove liveness at 0.85 percent.

Fraud in the Lending & Investment Industry

Altered documents lead the sector at 7.36 percent, ahead of replay and screenshots at 3.32 percent.

INDUSTRY FRAUD EXPOSURE

17.08%

of all verification requests in Lending & Investment were identity fraud in the first half of 2026

Identity fraud in lending and investment reached 17.08 percent of all verification requests in the first half of 2026. The Federal Reserve documents synthetic identity fraud as a structural weakness in lending, where fabricated and real data are combined into new identities that pass conventional onboarding and build credit histories before defaulting. It reports that 85 to 95 percent of applicants later identified as potential synthetic identities are not flagged by traditional fraud

models, so credit is extended to people who do not exist.¹² Shufti’s data shows altered documents leading the sector at 7.36 percent, the document artefact such applications rely on, followed by replay and screenshot attempts at 3.32 percent and AI and deepfake images at 2.99 percent, which is the highest reading for that instrument in any sector. Failure to prove liveness sits at 1.71 percent.

ATTACK TYPOLOGY BREAKDOWN




Figure 12. Lending & Investment: detection rate by attack instrument.

Fraud in the Telecom Industry

Altered documents lead the sector at 4.00 percent, ahead of replay and screenshots at 2.72 percent.

INDUSTRY FRAUD EXPOSURE

8.07%

of all verification requests in Telecom were identity fraud in the first half of 2026

Identity fraud in telecom reached 8.07 percent of all verification requests in the first half of 2026. The Federal Communications Commission describes SIM-swap fraud as a bad actor convincing a carrier to transfer a victim’s mobile service to a device in the attacker’s possession, without ever touching the physical phone. Once the number is moved, the attacker can intercept SMS one-time passcodes and take over the bank and email accounts tied to that number, a weakness the FCC attributes to

ATTACK TYPOLOGY BREAKDOWN




Figure 13. Telecom: detection rate by attack instrument.

breach-exposed personal data and weak subscriber-identity verification.¹³ Shufti’s telecom data is led by altered documents at 4.00 percent, the onboarding artefact behind fraudulent line activation, followed by replay and screenshot attempts at 2.72 percent and AI and deepfake images at 1.06 percent. Failure to prove liveness is effectively absent at 0.02 percent.

Fraud in the E-Commerce & Marketplaces Industry

Altered documents lead the sector at 5.60 percent, ahead of replay and screenshots at 3.75 percent.

INDUSTRY FRAUD EXPOSURE

13.29%

of all verification requests in E-Commerce & Marketplaces were identity fraud in the first half of 2026

Identity fraud in e-commerce and marketplaces reached 13.29 percent of all verification requests in the first half of 2026. The Federal Trade Commission ranks online shopping as the second most commonly reported fraud category in 2024, behind only imposter scams, and received more than 1.1 million identity-theft reports the same year. That places identity misuse and account fraud among the largest

consumer-exposure points in online commerce.¹⁴ Shufti’s data for the sector is led by altered documents at 5.60 percent, with replay and screenshot attempts also elevated at 3.75 percent, AI and deepfake images at 2.28 percent, and failure to prove liveness at 0.91 percent.

ATTACK TYPOLOGY BREAKDOWN




Figure 14. E-Commerce & Marketplaces: detection rate by attack instrument.

Fraud in the Forex Industry

Altered documents lead the sector at 4.88 percent, ahead of replay and screenshots at 3.39 percent.

INDUSTRY FRAUD EXPOSURE

17.18%

of all verification requests in Forex were identity fraud in the first half of 2026

Identity fraud in forex reached 17.18 percent of all verification requests in the first half of 2026. A joint CFTC and NASAA investor alert warns that retail foreign-exchange trading is a persistent target for fraudulent operators who solicit customers online but never place their funds in the market, diverting the money instead. The reliance on remote onboarding, and the difficulty of confirming that a dealer or account is legitimate, leaves retail participants highly exposed.¹⁵ For a forex business the threat is far from abstract, because the same forged documents, synthetic identities, and face swaps allow a single operator to open a cluster of

trading accounts under fabricated or stolen names and then exploit them for bonus abuse, multi-account trading, and the movement of value through the platform. Each account looks legitimate at the point of onboarding, so the chargebacks, the trading losses, and the regulatory scrutiny land on the broker long after the money has gone. Shufti’s data shows altered documents leading forex at 4.88 percent, followed by replay and screenshot attempts at 3.39 percent, AI and deepfake images at 2.80 percent, and failure to prove liveness at 0.47 percent.

ATTACK TYPOLOGY BREAKDOWN




Figure 15. Forex: detection rate by attack instrument.

Fraud in the Insurance Industry

Altered documents and replay and screenshots tie at the top, the only sector in the report where two instruments do.

INDUSTRY FRAUD EXPOSURE

15.05%

of all verification requests in Insurance were identity fraud in the first half of 2026

Identity fraud in insurance reached 15.05 percent of all verification requests in the first half of 2026. FATF’s risk-based guidance for the life insurance sector highlights non-face-to-face distribution channels, where the customer never physically meets an employee or intermediary, as an identity-verification weakness unless supported by other measures. It also flags arrangements where the policyholder, insured, and

beneficiary differ with no apparent relationship, and where beneficiary identity is often not verified until payout.¹⁶ Shufti’s data shows altered documents and replay and screenshot attempts level at the top of insurance, each at 4.18 percent, the only sector in the report where two instruments tie for the lead. AI and deepfake images follow at 1.74 percent and failure to prove liveness at 0.93 percent.

ATTACK TYPOLOGY BREAKDOWN



Figure 16. Insurance: detection rate by attack instrument.

Fraud in the Healthcare Industry

Altered documents lead the sector at 5.79 percent, ahead of replay and screenshots at 3.87 percent.

INDUSTRY FRAUD EXPOSURE

15.94%

of all verification requests in Healthcare were identity fraud in the first half of 2026

Identity fraud in healthcare reached 15.94 percent of all verification requests in the first half of 2026. The HHS Office of Inspector General describes medical identity theft as criminals stealing a patient’s name, Social Security number, or Medicare number to obtain care, acquire drugs, or submit fraudulent bills in the victim’s name. Because those identifiers are used at the point of care and billing, the fraud can

corrupt medical records and drain taxpayer-funded programmes.¹⁷ Shufti’s data shows altered documents leading healthcare at 5.79 percent, ahead of replay and screenshot attempts at 3.87 percent, AI and deepfake images at 2.63 percent, and failure to prove liveness at 0.88 percent.

ATTACK TYPOLOGY BREAKDOWN




Figure 17. Healthcare: detection rate by attack instrument.

Building Identity Resilience: A Multi-Layered Approach



Building Identity Resilience

Catching the forgery without catching the infrastructure only removes one disposable identity from a cheap supply

The recommendations resolve into one framework rather than a checklist. Detecting the presented artefact, whether a deepfake, a liveness failure, a document forgery, or a biometric spoof, is necessary and not sufficient on its own. Device fingerprinting and behavioural biometrics are what expose the network behind it.

- **Multi-layered verification:** document authentication, biometric matching, liveness, and deepfake detection read together at onboarding, so no single presented artefact is trusted alone.
- **Identity network intelligence:** device fingerprinting, IP and network analysis, and reuse and shared-attribute analysis across documents and devices.
- **Behavioural and continuous verification:** behavioural biometrics, with trust re-checked at login, recovery, and high-value events, so a network cannot quietly exploit an account once it has cleared onboarding.
- **Risk-based decisioning:** the evidence is weighted, clean attempts clear automatically, and edge cases route for review, so controls scale without adding friction for genuine users.

Continuous Identity Verification

An identity that cleared onboarding cleanly may still be the fourth of seventy

Identity assurance should continue past onboarding as an ongoing capability. The evidence in this report makes the case directly, because a coordinated network is defined by behaviour across time and across accounts, and a control that fires once at account opening can only ever see one moment of it. The signal that would reveal it arrives afterwards.

Trust should therefore be reassessed at login, account recovery, high-value transactions, beneficiary changes, privileged actions, and periodic KYC reviews, using the same network-intelligence signals that expose coordinated fraud at onboarding.

THE PRACTICAL SHIFT

From a pass or fail at the door to a running assessment

An assessment that can be revised as new evidence arrives, whether that evidence is a document reappearing under another name, a device resurfacing across unrelated accounts, or behaviour diverging from the customer it is supposed to belong to.

Conclusion

Identity fraud increasingly shows patterns of coordination

Digital trust is preserved instrument by instrument, and it is held only for as long as every instrument is held together.

Attackers reuse counterfeit documents, spoofed biometrics, identities, devices, and digital infrastructure across multiple verification attempts. Rather than relying on a single attack method, these networks exploit multiple layers of the identity ecosystem to create, replicate, and operate fraudulent identities across platforms.

Stopping fraud of this kind, whether it surfaces as bonus abuse, mule-account activity, multi-account trading, cross-border recruitment fraud, or other forms of identity misuse, asks more of identity verification than a single check at the door. It takes a layered mechanism that first defeats the artefact, detecting fake documents, catching deepfakes, and enforcing spoof and liveness detection, and then recognises the same actor across the signals a network cannot hide.

The response is therefore twofold, multi-layered verification at onboarding and continuous verification across the lifecycle, moving organisations from static checks to dynamic, risk-based assurance that catches both the individual fake and the coordinated network behind it.

The Identity Verification Solution

The capabilities that decide whether a network is caught or banked

Document verification, face verification, and liveness can all be present while the coordinated operation running many identities behind them goes unseen.

NETWORK AND FRAUD-RING DETECTION

- ☐ Can it identify potentially coordinated fraud activity by linking separate verification attempts through reused fraudulent documents, matching faces, shared devices, and controlled IP addresses?
- ☐ Can it perform account deduplication, so that one person opening several accounts under different identities is surfaced rather than cleared one attempt at a time?
- ☐ Can it detect activity where a single operation presents from more than one jurisdiction?

ATTACK DETECTION

- ☐ Can it detect injection attacks, including virtual cameras, camera-feed replacement, and video streams fed into the pipeline without ever reaching a physical sensor?
- ☐ Can it detect presentation attacks, covering replayed images, screenshots, masks, cut-outs, and other physical-material spoofs?
- ☐ Can it detect face swaps and synthetic composites at the biometric and liveness stage?
- ☐ Can it detect deepfake and AI-generated documents as distinct from conventional forgery?

FRAUD ANALYTICS

- ☐ Does it provide device fingerprinting and behavioural biometrics as live decision signals, not only as logged metadata?
- ☐ Are those signals correlated with document and biometric outcomes inside a single risk decision?
- ☐ Is trust reassessed after onboarding, at login, account recovery, and high-value events?

RESEARCH AND MODEL CURRENCY

- ☐ What continuous research and development programme supports deepfake detection, and how often are models retrained against new generative techniques?
- ☐ How quickly is a newly observed attack method reflected in production detection?
- ☐ Can performance be evidenced against recent attack samples rather than published benchmarks alone?

EVIDENCE TO LOOK FOR

Ask for the numbers, not the feature list

Detection rates by attack instrument, the false-positive rate that accompanies them, and a demonstration of how linked accounts are surfaced for review.

References

All figures marked [VERIFY] require confirmation against the primary source before publication.

1	Federal Bureau of Investigation and Internet Crime Complaint Center, North Korean IT Worker Threats to U.S. Businesses, PSA I-072325-4-PSA, 23 July 2025.	7	Financial Action Task Force, Virtual Assets Red Flag Indicators of Money Laundering and Terrorist Financing, 2020.
2	Financial Conduct Authority, Proceeds of Fraud: Detecting and Preventing Money Mules, 19 October 2023.	8	FinCEN, Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions (FIN-2024-Alert004), 2024.
3	Hong Kong Monetary Authority, AML Regtech: Network Analytics, May 2023.	9	FinCEN, Financial Trend Analysis: Identity-Related Suspicious Activity, 2021 Threats and Trends, 2024.
4	Shufti, Deepfake Identity Fraud Index Report 2026. Category shares of AI-enabled fraud, January to May 2026. [VERIFY: the 80.10 percent document deepfake share does not reconcile with the 11.9 percent published for 2025.]	10	UK Gambling Commission, The 2023 money laundering and terrorist financing risks within the British gambling industry, 2023.
5	United Nations Office on Drugs and Crime, Transnational Organized Crime Threat Assessment: Southeast Asia 2026.	11	Financial Action Task Force, Illicit Financial Flows from Cyber-Enabled Fraud, 2023.
6	UK Gambling Commission, Identity Theft and Fraud. More than 5,900 scans of identity and financial documents connected to fraudulent gambling accounts generating nearly £80,000. [URL to be inserted from the original source list.]		

12

Federal Reserve, Synthetic Identity Fraud in the U.S. Payment System, 2019.
[VERIFY: lending and credit focus only.]

13

Federal Communications Commission, Report and Order, Protecting Consumers from SIM-Swap and Port-Out Fraud (WC Docket No. 21-341), 2023.

14

Federal Trade Commission, Consumer Sentinel Network Data Book 2024, 2025.

15

CFTC and NASAA, Investor Alert: Foreign Exchange Currency (Forex) Fraud.
[VERIFY: alert is undated and its cited cases predate 2020.]

16

Financial Action Task Force, Guidance for a Risk-Based Approach for the Life Insurance Sector, 2018.

17

HHS Office of Inspector General, Fight Back! Against Medical Identity Theft and Medicare Fraud. [VERIFY: brochure is undated.]

18

Industry fraud rates in Section 5, and all network figures in Section 4 including the 70-identity ring and its device, address, and document counts, are drawn from Shufti production data for the first half of 2026 and are not externally published.



Digital trust is preserved
instrument by instrument,
and it is held only for as long
as **every instrument is held
together.**

Shufti detects the artefact and the network behind it, across document, biometric, liveness, device, and network signals, in one risk-based decision.

[Request a demo](#)

