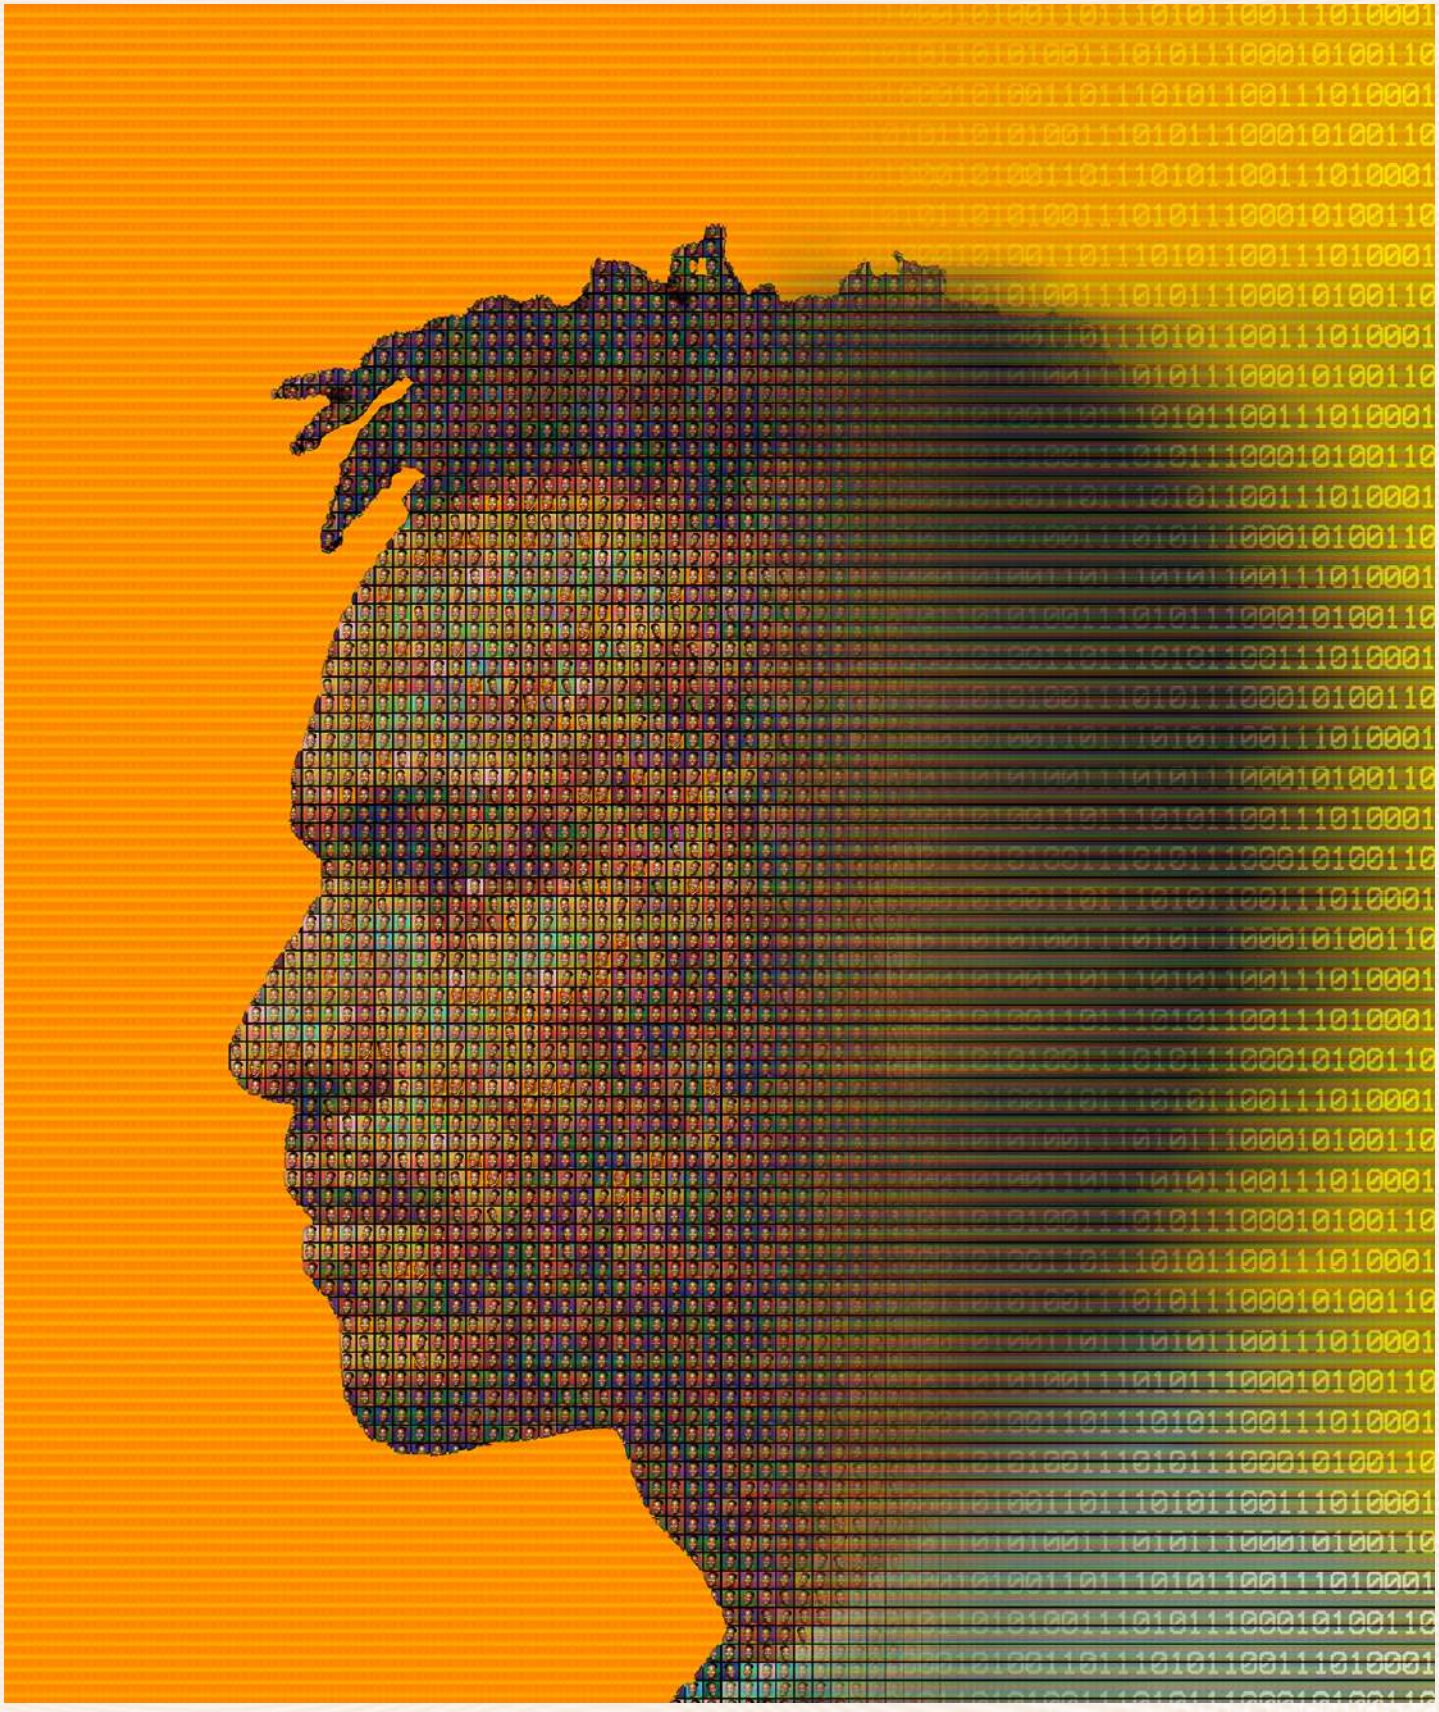

usesmileid.com

2026 Digital Identity Fraud in Africa Report

From Selfies to Signals: Identity Enters the Security Era

Smile ID

Table of Content

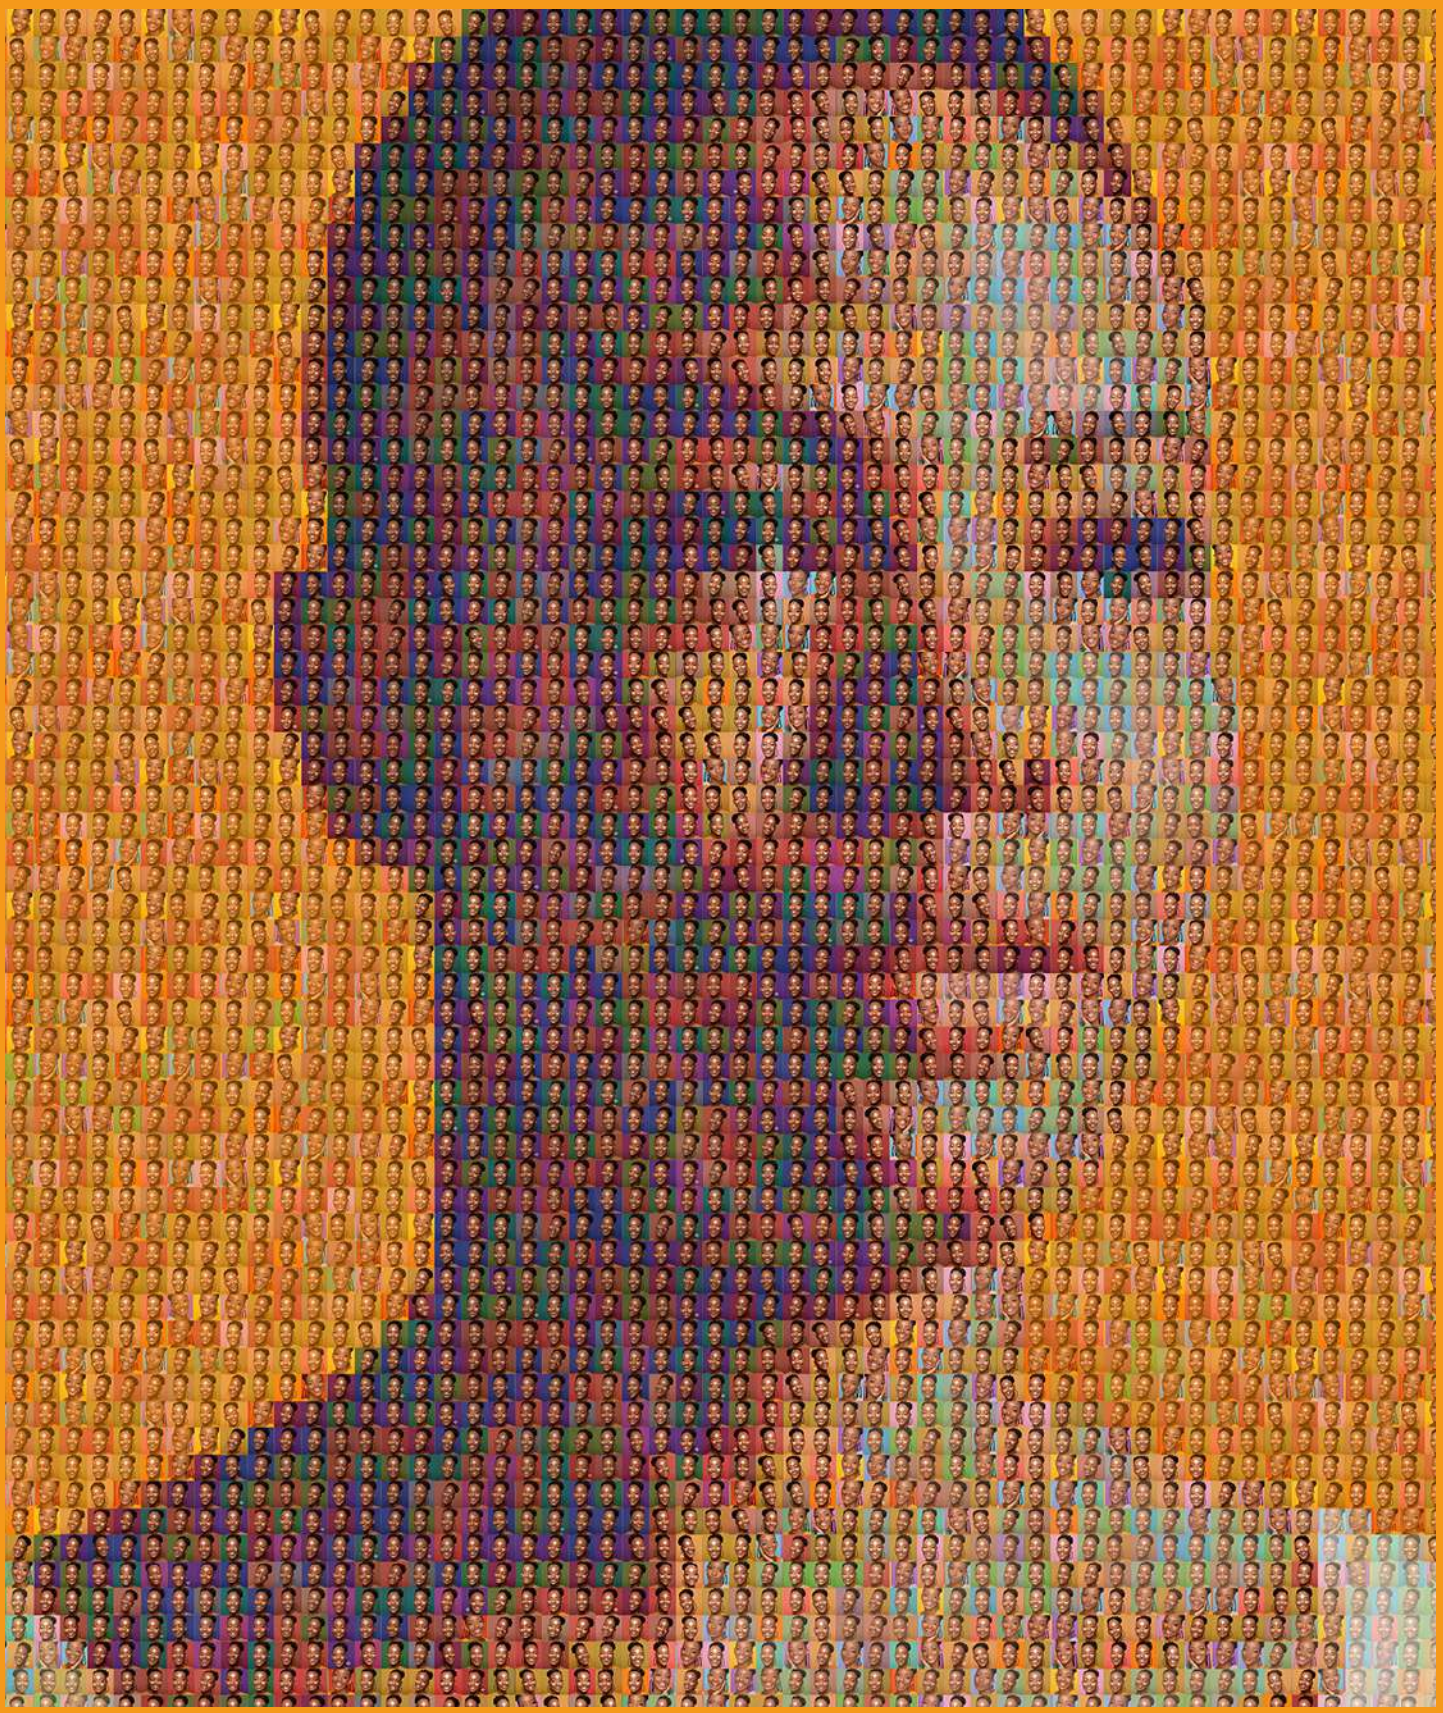

One face. Thousands of fraud attempts.

Foreword

Africa's digital economy is expanding faster than the security infrastructure that supports it. Over the past decade, the percentage of adults in Africa owning a financial account rose from 34% to nearly 60%, creating more than 200 million new accounts across the continent. That expansion has unlocked extraordinary opportunity, but it has also created a structural vulnerability: identity verification systems remain largely designed for a one-time checkpoint model, while fraud has evolved into continuous, industrial-scale operations that exploit security gaps across the full customer lifecycle.

In a single month in 2025, Smile ID traced more than 160,000 fraudulent verification attempts back to just 100 facial identities. Some of these faces appeared over 12,000 times across multiple platforms. In another case, attackers used the same identity to attempt more than 1,000 account registrations within 30 minutes. These patterns are not isolated incidents. They represent coordinated fraud networks operating at scale—reusing stolen identities, automating attacks, and targeting the moments where value concentrates: login flows, password resets, device changes, and high-value withdrawals.

At Smile ID, our mission extends beyond verification. We are helping build the trust infrastructure—systems that verify identity, detect fraud, and protect accounts—that Africa's digital economy now urgently requires.

Using insights from 200 million identity checks across Africa in 2025, this report reveals how fraud tactics are changing, where legacy controls fall short, and how digital platforms can reduce fraud without losing legitimate customers as they enter 2026.

~200M
New financial accounts created across Africa in the last decade.

160k+
Fraudulent verification attempts tracked by Smile ID in a single month in 2025.

200M+
Identity checks Smile ID conducted across Africa in 2025.

The most consequential fraud attacks today are targeted account takeovers (ATOs)—not fake IDs or isolated spoofs, but coordinated operations that compromise the capture pipeline, reuse real identities at scale, and exploit moments after approval when controls are lighter through highly scalable AI-powered tooling.

From Selfies to Signals: Identity Enters the Security Era

For years, identity verification in Africa operated as a checkpoint: capture a selfie, match a face, verify a document. If the images looked legitimate and the biometrics aligned, the system granted access. That model assumed fraud was rare, isolated, and detectable through visual inspection. It worked when attackers were mostly individuals with limited resources, attempting to bypass onboarding with fake documents or low-quality spoofs.

By 2025, that assumption had collapsed. Fraud is no longer dominated by one-off deception. It has become repeatable, automated, and coordinated—targeting specific high-value accounts not at onboarding, but during login, account recovery, re-authentication, and withdrawal. Attackers operate as syndicates with repositories of facial and identity data, running scalable AI-powered attacks on a daily basis, often in bursts across multiple financial service providers. They return repeatedly. They hijack or purchase verified accounts. They reuse legitimate identities. And they move across platforms that lack shared risk intelligence.

This shift changes what matters. **The most consequential fraud attacks today are targeted account takeovers (ATOs)—not fake IDs or isolated spoofs, but coordinated operations that compromise the capture pipeline, reuse real identities at scale, and exploit moments after approval when controls are lighter.** These attacks take place later in the customer journey and are perpetrated by professional fraudsters, often in collusion with insiders, who leverage repositories of facial and identity data through highly scalable AI-powered tooling. They iterate attacks daily, often in bursts across multiple financial service providers. The challenge is no longer verifying identity once. It is protecting identity continuously across the full customer lifecycle against adversaries who treat fraud as infrastructure. Effective defence requires connecting the dots on fraudulent activity across attempts, clients, and time, so patterns are identified early and preventive actions taken. Understanding why this shift occurred requires examining the economics that enable it.

The Marginal Cost of Fraud

Generative AI has fundamentally altered the cost structure of attempting fraud. High-quality synthetic documents, deepfake imagery, and automated biometric manipulation are no longer expensive or rare. What once required specialist skills and significant time can now be produced cheaply, repeatedly, and at scale. When fraud is cheap, attackers don't need to succeed on the first attempt. They test systems continuously and iterate until they break through.

Defences built for rare or one-off attacks cannot withstand constant, automated pressure. Attackers can afford to reuse the same identity assets across hundreds or thousands of attempts because the marginal cost of each try approaches zero. The return on investment for fraud infrastructure—face repositories, injection tooling, emulator farms—increases as that infrastructure is amortized across more attacks, more platforms, and longer campaigns.

Any control assuming scarcity of attempts, identities, or attacker capability, will be systematically overwhelmed. Fraud defence must now assume abundance and use networked intelligence to spot patterns and turn the volume generated by fraudsters' attacks against them.

Three Fundamental Changes Reshaping Fraud in 2025

1. Fraud Moved from Onboarding to the Full Lifecycle

In 2025, authentication-related fraud attempts were five times more common than onboarding fraud. Attackers are no longer focused on creating new accounts. They target already verified users at the moments that unlock value: login and re-authentication, password resets, device changes, and high-risk transactions like withdrawals or limit increases. These flows are often designed for user convenience, which makes them structurally weaker than onboarding. They rely on lighter controls such as SMS OTPs, email links, security questions, that can be bypassed through SIM swaps, social engineering, or insider assistance.

This pattern was especially visible in West African retail banking, where potential fraud attempts rose approximately 50% year-over-year, driven primarily by verification volumes in authentication and account recovery flows. The accounts being attacked had already passed KYC. Most fraud did not happen at the perimeter. It happened inside trusted systems, exploiting gaps in post-approval identity assurance.

Controls designed only for onboarding are structurally mismatched to this threat. Effective defence now requires verifying identity not once, but continuously treating each high-risk interaction as a moment that demands proportionate proof of presence and ownership.

2. Fraud Became Repeatable, Automated, and Cross-Platform

Modern fraud operates as coordinated campaigns that reuse the same identity assets, namely faces, devices, documents, and behavioural patterns, at extreme scale across many platforms.

Individual attempts can appear legitimate in isolation. A selfie matches. A document looks clean. The device seems normal. But when viewed across sessions and platforms, these attempts follow a repeatable playbook: attackers exploit the weakest entry point, reuse compromised or purchased identities at high-value moments, and move quickly before review systems can respond. This is networked fraud, and it cannot be stopped with isolated defences.

In 2025 Smile Secure, Smile ID's biometric deduplication capability, detected 71% more duplicate fraud attempts than the combined total of 2023 and 2024. This growth reflects not just increased attack volume, but increased sophistication: syndicates now operate supply chains for identity assets, "aging" accounts through dormancy or low-risk activity before activating them for fraud or money laundering.

Fraud detection must therefore recognize repetition, reuse, and coordination across time, sessions, and platforms, not just evaluate single transactions in isolation.

3. Fraud Moved from Visual Deception to Pipeline Compromise

As fraud becomes automated and repeatable, a selfie is no longer a sufficient control. Attackers can reuse real identities, inject compromised inputs, or run attempts from manipulated environments. Even when an image looks valid, the surrounding system may not be trustworthy.

In 2025, nearly 90% of verifications rejected for suspected fraud were caught using mobile SDK integrations, up from 15% in 2023 and 65% in 2024. This is because SDKs capture additional on-device signals that API-only flows cannot see: image capture integrity, device status, user behaviour, and environmental context. API-based verification systems see only the final output; a selfie, a document, with limited visibility into how that media was produced. If the camera feed has been replaced with a virtual camera, if the app environment has been tampered with, or if the device is an emulator running automated scripts, the verification data is compromised before any analysis begins.

Smile ID's clustering analysis flagged more than 100,000 injection-style attempts per month in 2025, linked to emulator farms, virtual cameras, or tampered capture setups. These attempts often look visually normal when viewed individually. The fraud is revealed not through the image itself, but through metadata inconsistencies, abnormal capture timing, conflicting hardware and software fingerprints, and patterns of reuse across coordinated networks.

Fraud defence has therefore expanded from visual proof to system signals. Trust can no longer be inferred from images alone. It must be established through the integrity of the systems that produce them: secure capture environments, device-level validation, and networked intelligence that connects attempts over time.

This does not mean abandoning biometrics or documents. It means placing them within a broader defence system: trusted capture, layered controls across the users' lifecycle, device intelligence and patterns across a network.

90%

90% of rejections for suspected fraud are caught using mobile SDKs.

100k+

Injection-style attempts flagged per month in 2025 by Smile ID's clustering analysis.

What This Means for Defence in 2026

Together, these three changes create a clear operational reality: fraud attempts are now continuous throughout the customer journey, coordinated across platforms, and capable of bypassing controls that only assess what is submitted without validating how it was created.

Overall verifications rejected by Smile ID for suspected fraud declined from 25% in 2024 to 22% in 2025. This is driven by increasingly targeted fraud, rather than reduction in fraud attempts. Petty onboarding fraud attempts declined while attackers focused resources on sophisticated account takeovers of high-value accounts. The absolute number of fraud attempts continued to rise as verification volumes grew, but the nature of those attempts shifted toward precision rather than volume.

Defending against this requires moving beyond one-time KYC to continuous risk assessment across the customer lifecycle. It requires systems that can detect identity reuse across sessions and platforms, harden authentication at high-risk moments, and validate capture integrity at the source. Identity is no longer about KYC compliance. It has become security infrastructure.

The remainder of this report examines the evidence behind these changes in detail: the specific fraud techniques being deployed (Section 1), how those techniques manifest differently across African regions based on infrastructure and regulatory context (Section 2), and the practical defence strategies institutions must prioritise to protect identity across the full lifecycle (Section 3).

Methodology and Scope

This report analyses over 200 million identity verification checks conducted across Africa in 2025, drawn from a cumulative dataset exceeding 400 million checks processed across the Smile ID network since 2019. The data spans 37 industries across 35+ African countries.

We employed descriptive analysis, statistical modelling, and pattern-based review across biometric, document, device, and behavioural signals. The analysis covers the full identity lifecycle—onboarding, authentication, and high-risk account events—examining how fraud manifests at different stages of trust.

Quantitative findings were cross-referenced with qualitative insights from fraud, risk, and compliance professionals across multiple African markets, alongside operational experience from Smile ID's internal risk and engineering teams.

All data is aggregated and fully anonymised. Any document samples or imagery featured are illustrative only and do not represent real individuals. Key terms are defined in the Glossary at the end of this report.

Imagery

The ID documents featured in this report were digitally created using vector and raster graphic editing software to accurately resemble actual documents. These examples are for illustrative purposes only and must not be used to create fraudulent identities. All human imagery in this report consists of stock photographs digitally enhanced and refined using technology.

Highlights from the Data in 2025

5x

More Fraud at Authentication than Onboarding

Authentication is now the primary battleground. Authentication-related fraud attempts are 5× more common than at onboarding.

160k+

Fraud attempts traced to just 100 faces

Fraud is repeatable, automated, and cross-platform. Attackers reuse the same identities at scale: in one month, 100 real faces were reused across 160,000+ attempts.

90%

Fraud caught via SDK integration

SDK capture has become a signal advantage. Nearly 90% of rejections due to suspected fraud are caught using mobile SDK integrations. Mobile SDKs can capture additional on-device signals that APIs don't see.

1k+

Attempts from one identity in 30 minutes

Generative AI Has Collapsed the Economics of Fraud. High-quality synthetic documents, deepfakes and images are no longer expensive to produce or iterate on. Smile ID detected relationships between images that were tried 12,000+ times across platforms. In another case, one synthetic identity was re-tried over 1,000 times within 30 minutes.

100k+

Injection-style attempts per month

Metadata and device-level signals are now a critical detection layer. Leveraging cluster data (shared devices, networks or capture patterns), Smile ID detects 100,000+ injection-style attempts per month that appear linked to emulators, virtual cameras, or tampered capture.

How Fraud Has Changed: From Fake Accounts to Account Takeover

This section breaks down the mechanisms shaping identity fraud in 2026, discussing common fraud types and techniques.

01

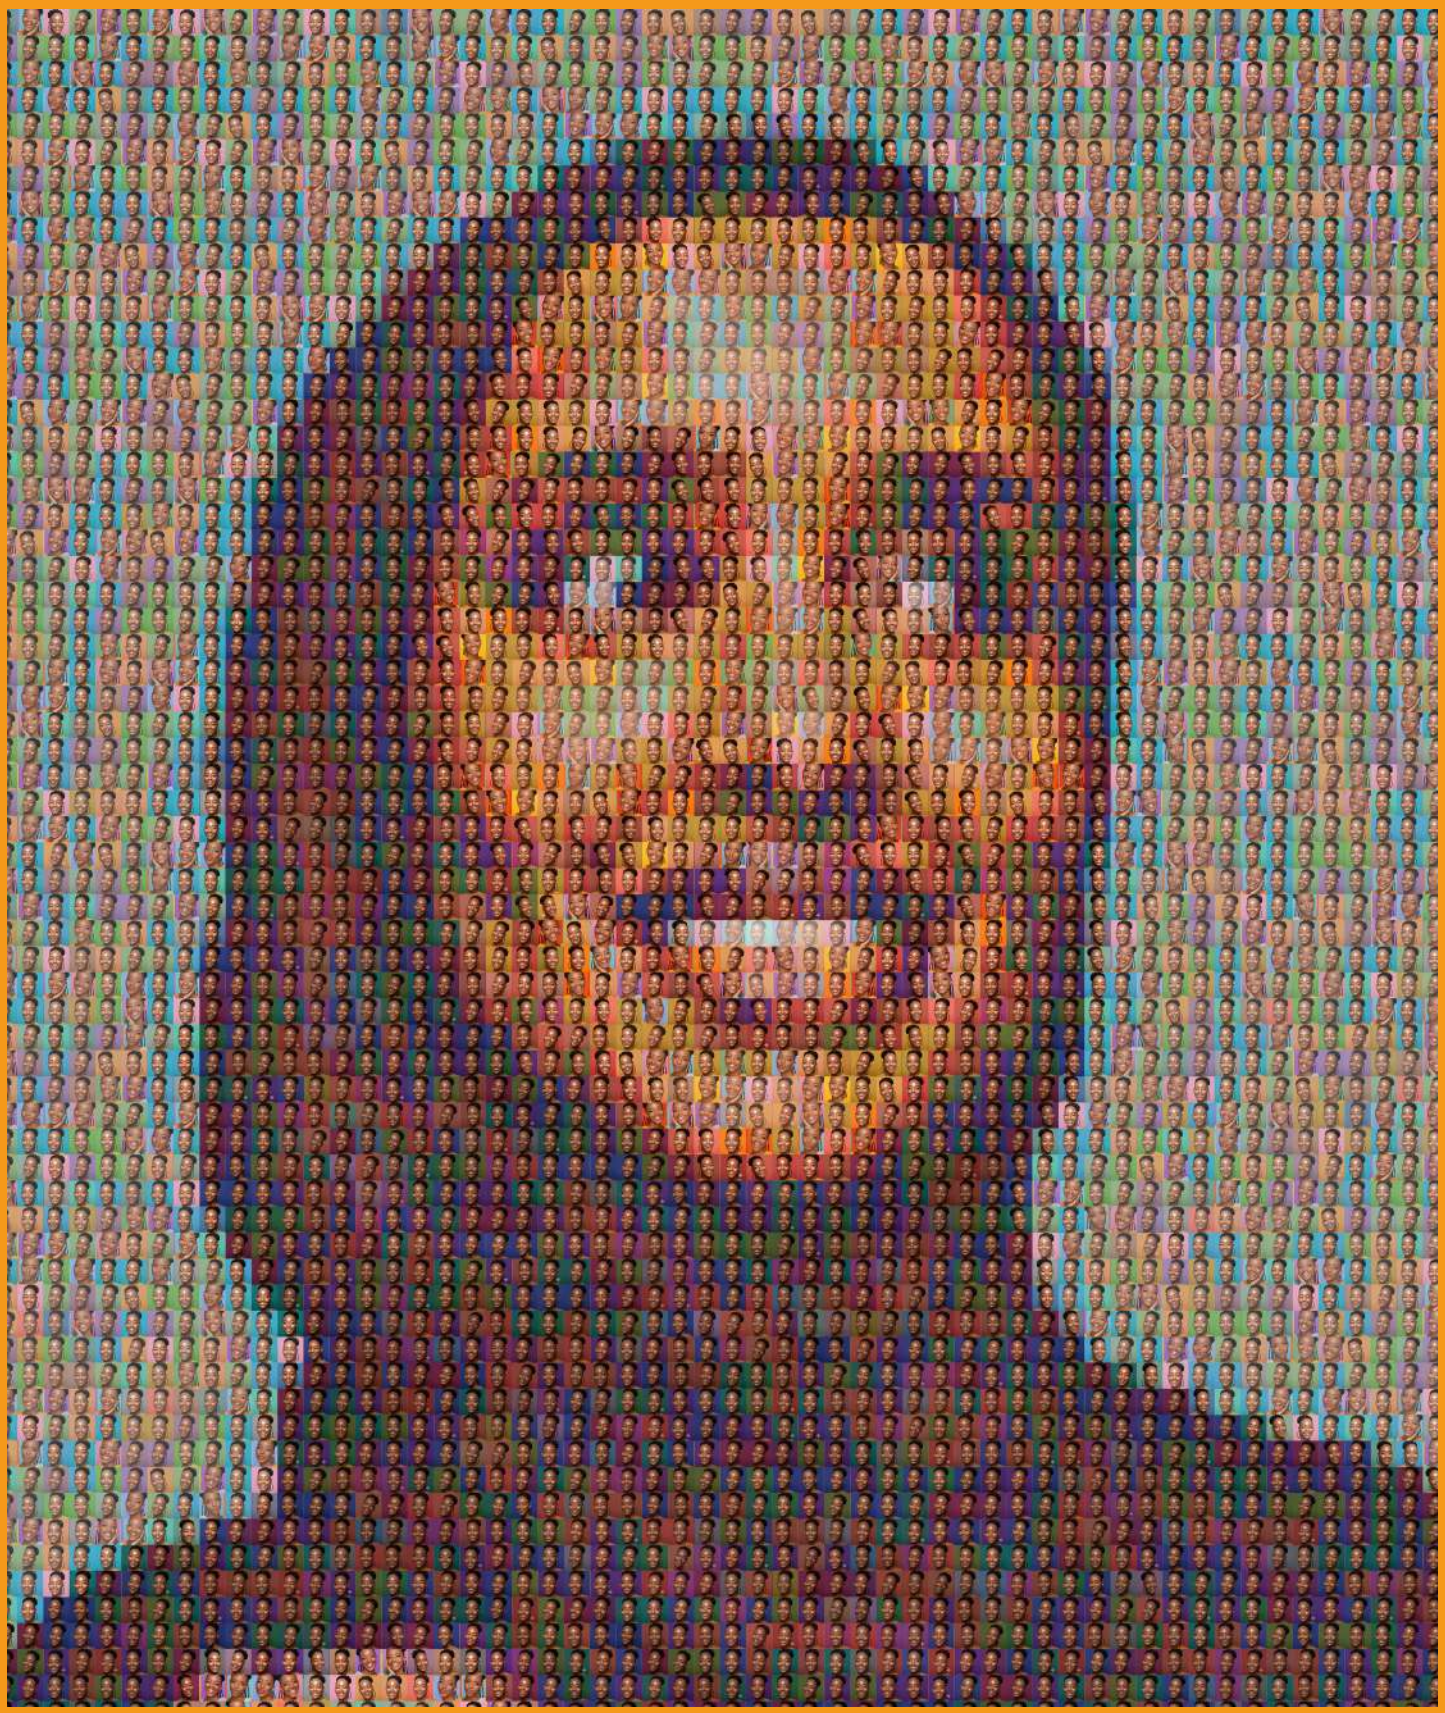

One face. Thousands of fraud attempts.

1.1 Facial Biometric Fraud

This section covers biometric fraud in two parts. First, it summarises the key indicators that show up at scale during facial verification and trigger rejection or investigation. Second, it highlights a subset of cases, where the attempts were analyst-reviewed for confirmation, and the techniques were classified.

These views show the prevalence of different attack methods as a % of all biometric fraud detected and what they have in common.

Prevalent Types of Biometric Fraud Attempts in 2025

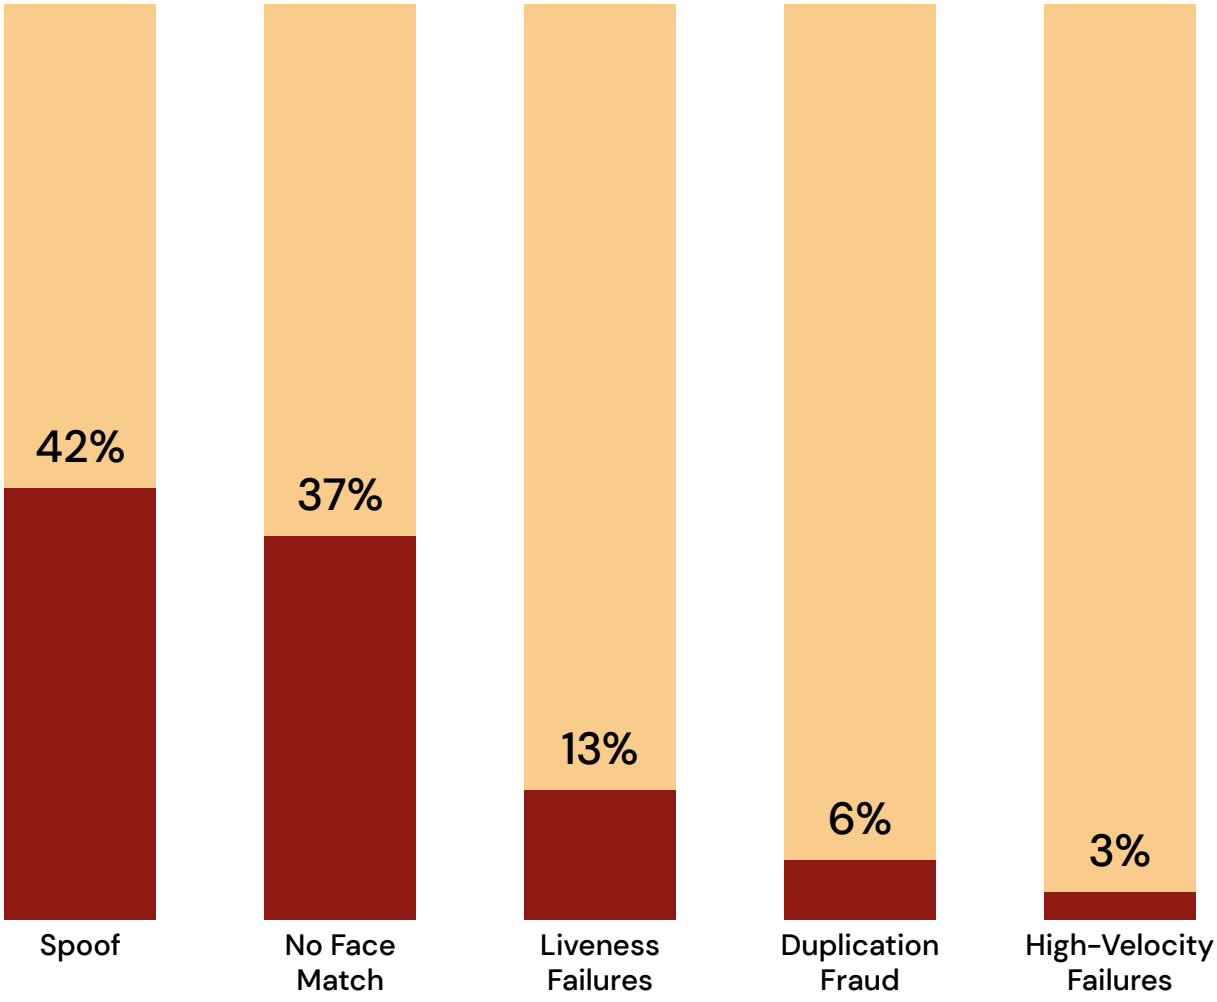

Within the set of attacks classified as biometric fraud, the most common categories include:

- **Spoof attempts:** Attempts to pass verification using photos, videos, or manipulated facial media.
- **No face match:** The submitted selfie does not match the claimed identity or document image.
- **Liveness failures:** Detection of replayed content, non-human behaviour, or synthetic motion.
- **Duplication:** The same face appearing across multiple accounts or verification sessions.
- **High-velocity failures:** Multiple failed attempts in a short period, indicating automation or systematic probing rather than normal user error.

While any individual category may result in an automatic rejection, when analysed together over time, we are able to see clusters—repeated attacks using similar combinations of techniques across many sessions or devices.

Confirmed Biometric Fraud Techniques

Smile ID analysts reviewed and classified thousands of 2025 fraud cases using the biometric media plus supporting signals like device posture and metadata.

Three technique patterns show up most often in confirmed biometric fraud:

Breakdown of Known Biometric Fraud in 2025

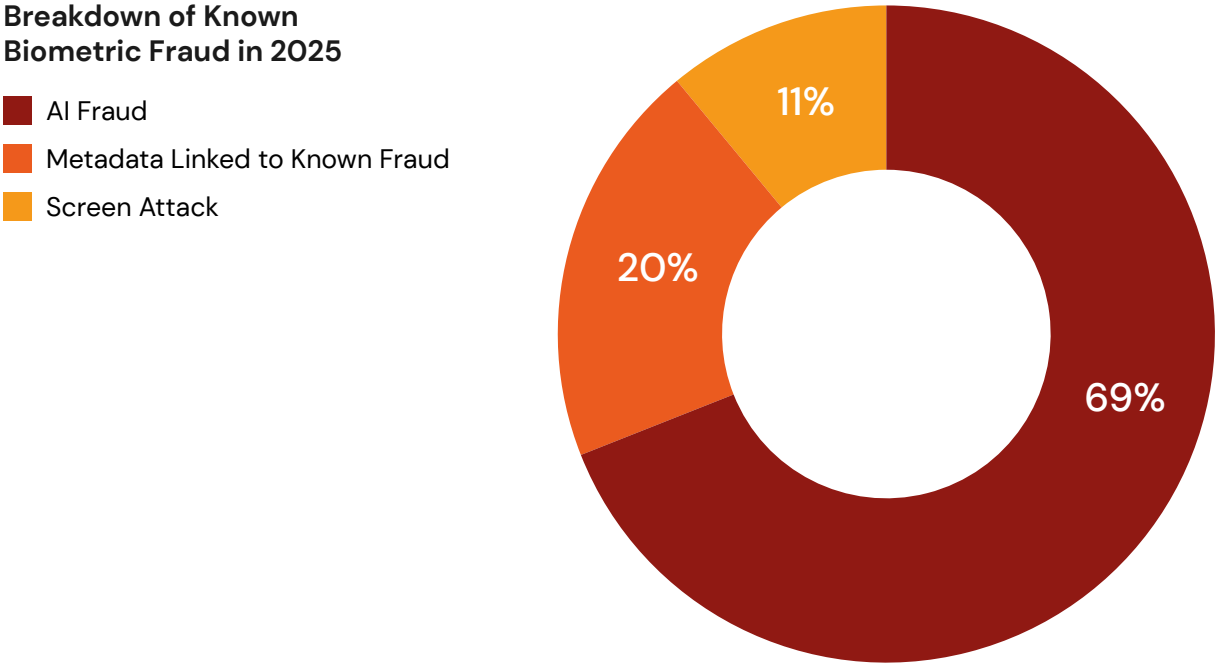

1.1.1 AI-Generated Manipulation at Scale (69% of confirmed cases)

This is the most common pattern. Attackers use AI to create or alter face media so it can pass verification. Examples include synthetic faces, deepfake video made from still images, face swaps, and hybrid overlays that mix real footage with AI-generated facial elements. The goal is to produce media that can pass liveness and similarity checks.

Early-stage attacks

2019-2020

Hidden cameras capturing videos of legitimate users to bypass biometric checks.

Scaling abuse

2020–2021

The same real face reused across multiple accounts to exploit platform-level silos.

Physical deception

2021–2022

Cardboard cut-outs, masks, or printed images used to simulate presence.

Transitional manipulation

2023

Basic facial edits and overlays designed to bypass visual inspection.

Face swaps & AI-assisted fraud

2024

Synthetic techniques paired with identity reuse and automation.

Identity Hijacking

2025

Real-time grafting of identities to bypass modern security measures.

1.1.2. Replay and Injection Attacks Bypassing Controls (11% of confirmed cases)

These attacks reuse previously captured photos or videos and feed them into the verification flow through screens, virtual cameras, or injected media streams.

Common methods include photo-of-a-screen, video replay, screen-on-screen attempts, and replay routed through virtual cameras. They are less advanced than AI manipulation, but can still work if capture quality is weak and advanced liveness checks are not used.

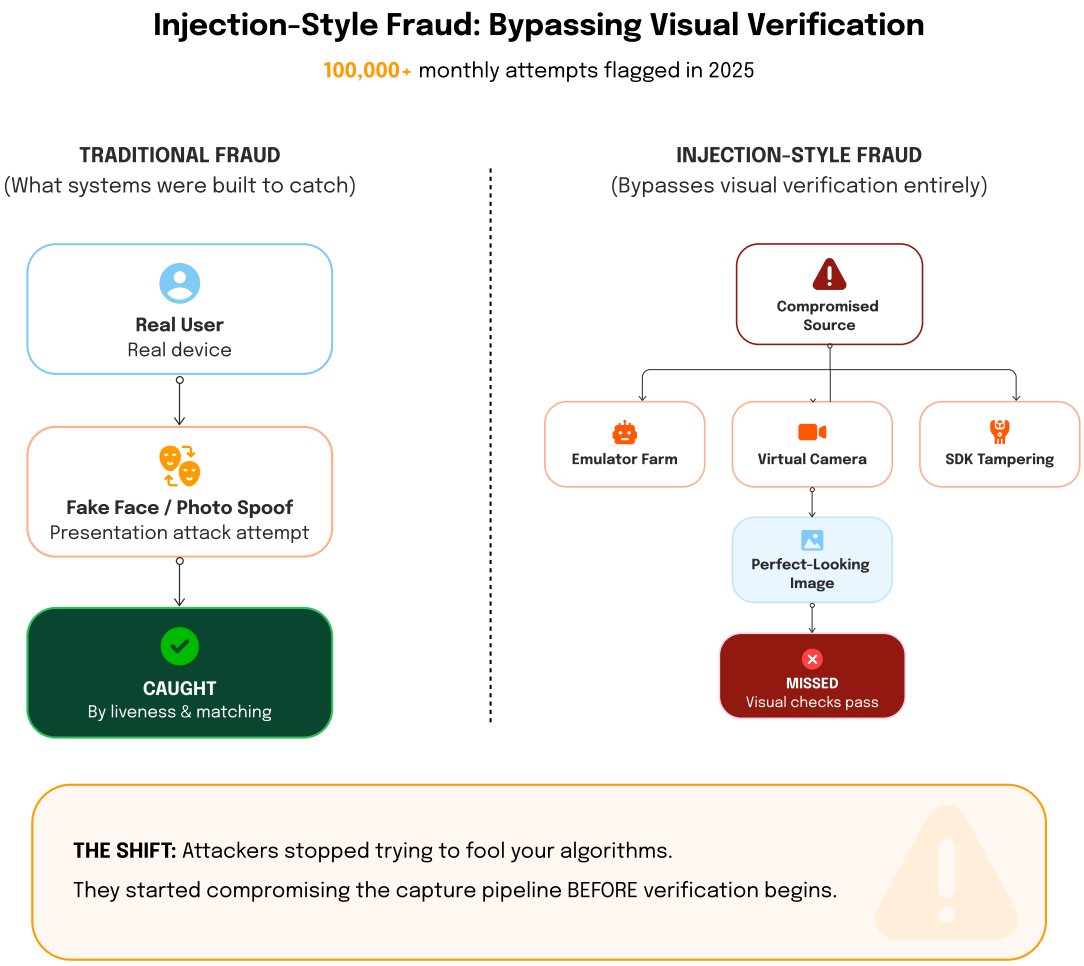

1.1.3. Metadata Linked to Known Fraud (20% of confirmed cases)

In these cases, the face media can look legitimate, but fraud is confirmed through non-visual linkage signals, such as:

- The same face embedding reused across multiple accounts or platforms.
- Devices or environments already linked to confirmed fraud.
- Metadata information consistent with fraud automation, emulators, or injection attacks.

In real attacks, these patterns often overlap—an AI-manipulated face may be delivered via replay/injection and tied to reused devices or identities.

Biometric Fraud is focused on Authentication

Biometric fraud now happens mostly during authentication and other post-onboarding flows. As mobile banking grows, fraudsters target already verified accounts, where they can access funds quickly, use higher limits, and cash out faster than they can through new-accounts.

Authentication vs Registration Fraud Rate in 2025

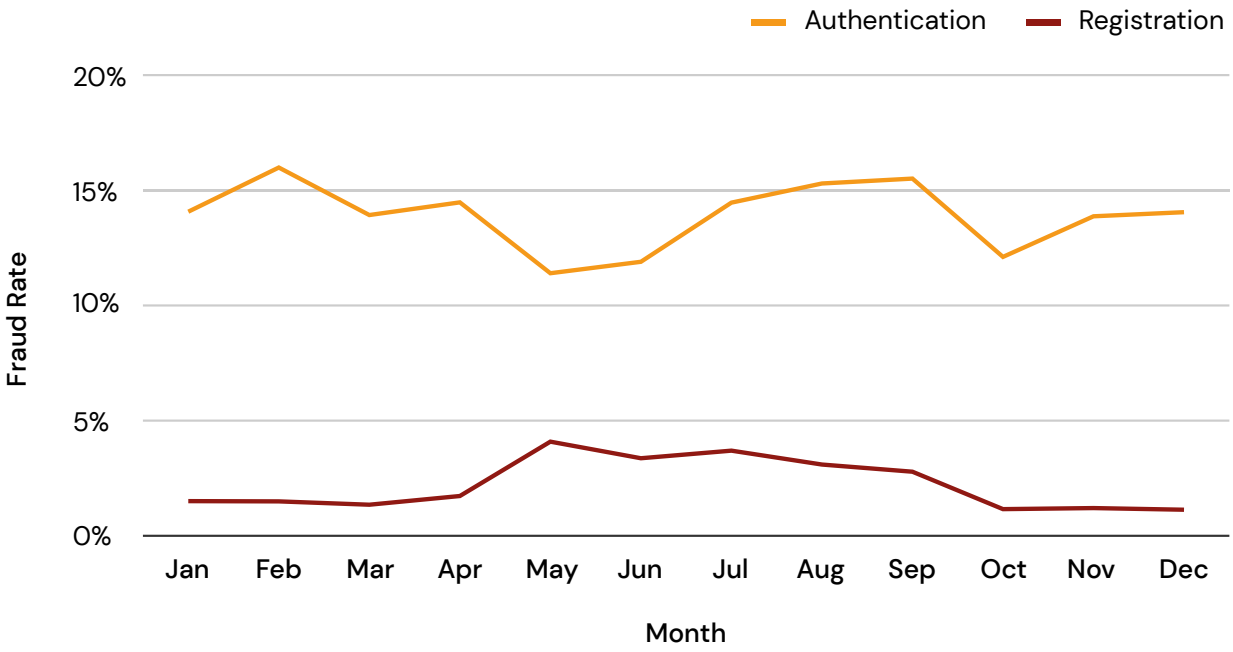

Are you human?

One of these faces is real. The others were generated by AI in seconds.

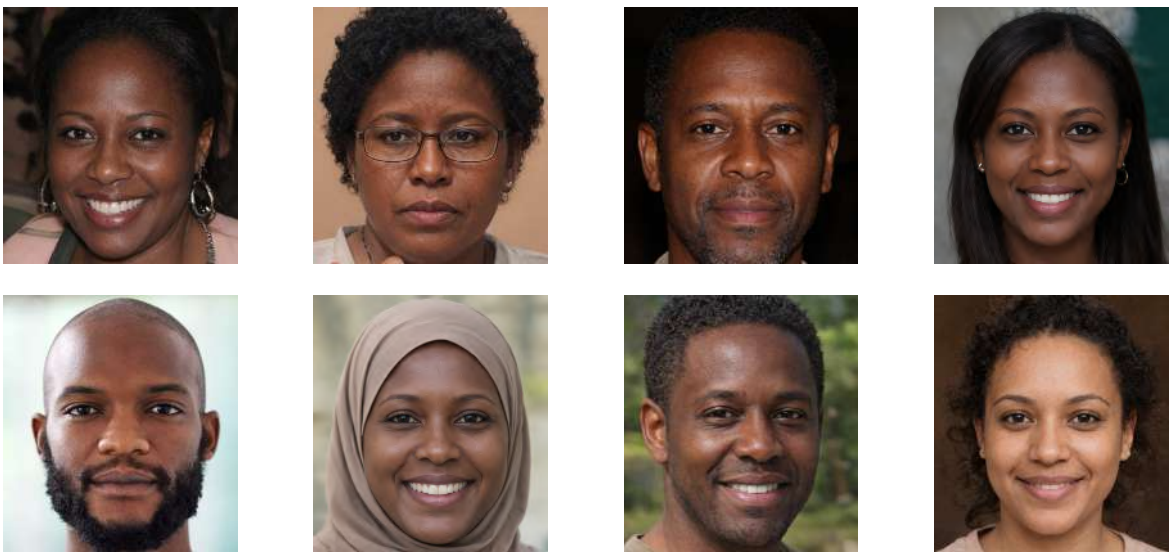

1.2 Injection Attacks and Device Manipulation: Attacking the Security Layer

Biometric fraud deceives what a system sees; injection attacks compromise how that biometric is captured and delivered. Instead of presenting a real face, attackers replace or manipulate the capture stream so the system receives a seemingly valid, live input without any genuine camera capture. The most sophisticated attacks now target the capture pipeline, not just the verification model.

Why Injection Attacks Matter Now

Injection-style attacks were historically less common in Africa. In 2025, Smile ID observed a shift toward infrastructure-level manipulation.

Techniques often associated with mature fraud environments; emulator farms (large networks of simulated devices used to automate and scale attacks); virtual camera feeds (spoofed or pre-recorded video streams designed to bypass liveness and face checks); and SDK tampering (modifying or manipulating a verification app's software components to disable protections or forge signals); are now being deployed at scale against African fintechs and banks.

This shift changes the defensive problem: when capture is compromised, visual plausibility is no longer a reliable control. A submission can look perfect and still be synthetic.

How Fraud Bypasses Systems Without "Beating the Models"

Injection attacks succeed where systems only verify what is submitted, but not how it was captured or produced. We have observed three common attack methods across the Smile ID network:

- **Virtual camera injection:** The physical camera is replaced with a software source, feeding pre-recorded or generated media into the flow. The output appears unusually clean and consistent — a signal that only surfaces when capture integrity is examined.
- **Emulators and scripted automation:** Fraud syndicates run hundreds of verification attempts in parallel using emulated devices, rotating identifiers and network attributes automatically. This turns biometric abuse into an assembly line rather than a one-off attempt.
- **SDK and payload manipulation:** Attackers tamper with mobile SDKs or intercept data in transit, injecting manipulated media and forging sensor metadata to mimic legitimate hardware during API submission.

Injection can happen at two points—during capture and during transfer. Effective defence must, therefore, validate the full pipeline, not just the final selfie.

What the Data Reveals About Scale & Automation

Smile ID's device and metadata checks flagged 480,000+ verifications with signs linked to injected or virtualised capture environments—such as conflicting hardware fingerprints, abnormal capture timing, and indicators of emulation or SDK tampering.

These signals become strongest when connected over time, revealing coordinated networks built to run fraud repeatedly and quietly.

480k+

Verifications flagged for signs of injection or virtualised capture environments.

1.3 Document Fraud

Generative AI has made document fraud harder to spot. Obvious screen and print tricks are giving way to high-quality digital forgeries that can look cleaner than real IDs captured in everyday conditions. This flips the old assumption: a sharp, "perfect" document image is no longer a sign that it's genuine.

Document Fraud as an Entry Point, Not the End Goal

In 2025, Smile ID saw a 250% increase in high-fidelity forgery signals (portrait anomalies) and a decline in older screen-based manipulation. The pattern is consistent with the broader shift across the report: fraud is becoming less noisy, less obvious, and more technical.

The chart below groups document fraud into a number of recurring manipulation techniques.

Breakdown of Document Manipulation Techniques in 2025

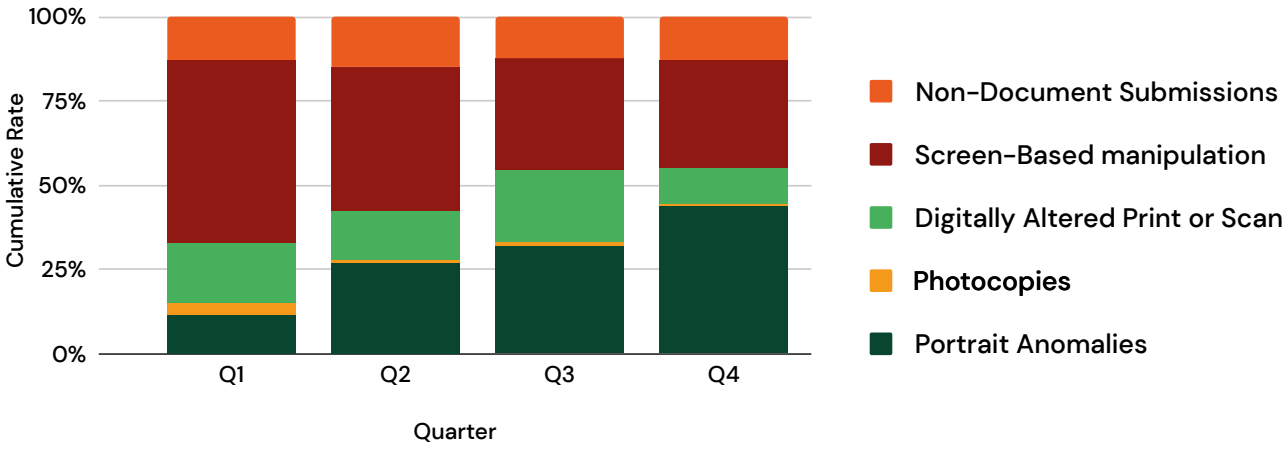

- **Non-document submissions:** Images that contain no valid ID (e.g., blank pages or unrelated objects), often used to probe system weaknesses.
- **Screen-based manipulation:** Digital images of identity documents displayed on a screen and re-captured, rather than photographed as a physical document.
- **Digitally altered prints or scans:** Documents that have been edited digitally and then printed or scanned to mimic authentic IDs, masking signs of manipulation.
- **Photocopies:** Copies of original documents that obscure security features such as holograms, textures, or watermarks.
- **Portrait anomalies:** Manipulation edits focused on the portrait area to realistically bridge the link between the ID and the presenter.

The Rise of High-Fidelity Document Fraud (Portrait Anomalies)

Portrait anomalies significantly increased across the year. Instead of editing the whole document, attackers often change only the portrait, swapping faces, adjusting alignment, or inserting synthetic facial images, while keeping the rest of the ID looking normal.

These edits are rarely obvious to humans. They are typically detected through cues like:

Lighting that doesn't match between the portrait and the rest of the document.

Unnatural face alignment or proportions.

AI Generation artefacts that don't match the surrounding document features.

Captures usually include blur, glare, uneven lighting, and background noise. Manipulated or AI-generated documents often look too uniform—clean edges, consistent lighting, and overly precise alignment. That's why detection increasingly relies on forensics and context, such as:

- Texture inconsistencies that do not match physical PVC or paper materials
- Digital artefacts suggesting an image was rendered or composited rather than captured
- Contextual mismatches in metadata that conflict with expected camera and device behaviour

Document fraud also rarely happens alone. A high-quality fake ID is often used to get through the door long enough for biometric reuse, injection, account takeover, or laundering to follow.

1.4 Duplication Fraud: The Signature of Fraud Syndicates

In 2025, duplication was one of the clearest signs of syndicate activity in Smile ID data. It shows up as the same face, device, or behaviour being reused across many accounts and platforms. What looks like isolated fraud to one business is often a coordinated operation against multiple businesses at scale.

Reuse Is the Clearest Signal

Smile ID saw this reuse model clearly in 2025:

- **The Hundred-Face Syndicate:** Smile ID traced back over 160,000 verification attempts across our network to a fraud syndicate using just 100 different faces.
- **Extreme Persistence:** A single "seed face" attempted to attack a partner platform over 12,000 times in a single month.
- **The "30-Minute Blitz":** One identity attempted 1,000+ onboardings on a platform in 30 minutes.

The defining feature is not how convincing one attempt looks—it's how aggressively the same assets are reused. Platforms that only see their own traffic often detect this pattern only after losses occur.

Duplicate fraud attempts more than doubled year over year—and nearly tripled the combined 2023–2024 total.

Duplicate Fraud Attempts caught by Smile Secure

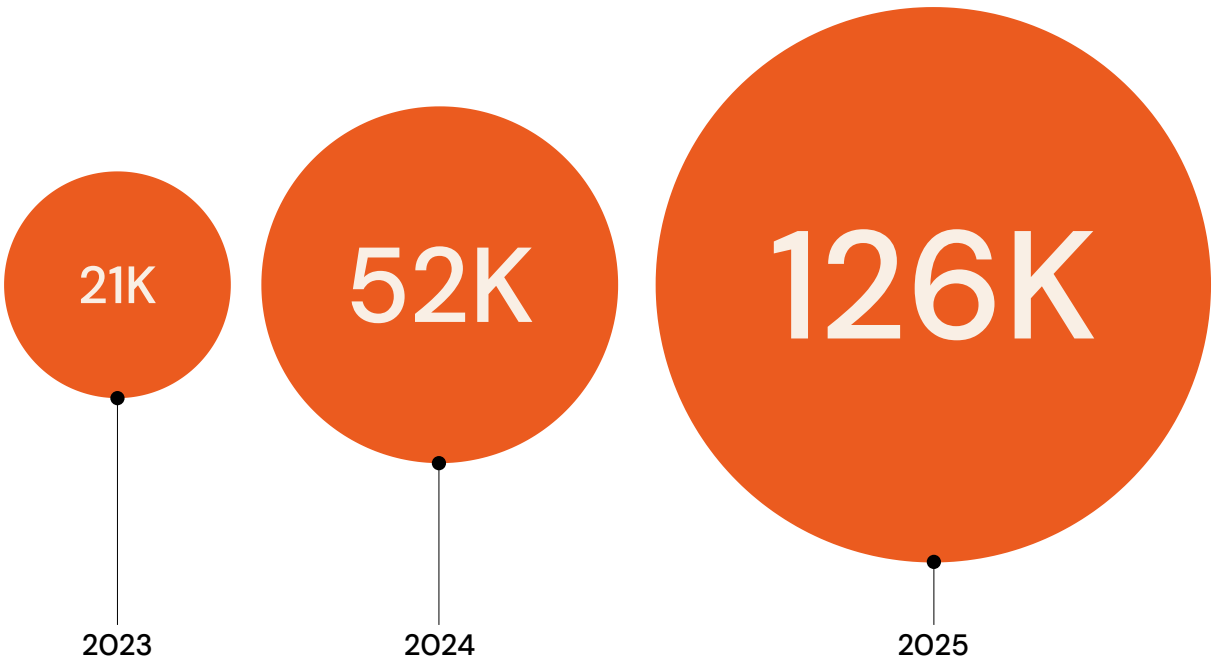

Implication for 2026

Duplication cannot be solved through tighter onboarding alone. Institutions must prioritise identity systems that:

- Detect reuse across time and sessions, not just single verification events
- Link biometric, device, and behavioural signals into repeatable patterns
- Support real-time intelligence sharing, where permitted under local data protection rules, to counter cross-platform syndicates

Fraud syndicates operate as ecosystems. Fraud defence must do the same.

One face. Thousands of fraud attempts.

1.5 Identity Farming & Money Laundering

While duplication is a consequence of syndicated fraud, identity farming is the supply chain behind it.

Identity farming has moved from a "rising tactic" to a mature operating model. Instead of creating or taking over accounts and cashing out immediately, fraud groups now build a supply of accounts over time, let them "age," and then use them in short, high-impact bursts.

From Single Accounts to ID Supply Chains

Modern identity farming works because the accounts look normal until they are activated. Common inputs include:

- Harvested or purchased real identities, often from individuals with limited digital footprints.
- Agent- or insider-assisted onboarding, where accounts are created or verified on someone else's behalf.
- Dormancy and low-risk activity, where accounts sit idle or perform small "normal" transactions to build history.

By the time these accounts are activated, they look legitimate, trusted, and low-risk—ideal vehicles for fraud or laundering.

How Identity Farming Enables Money Laundering

Once a pool of "trusted" accounts exists, funds can be moved quickly and at scale—through fintech apps, wallets, crypto exchanges, or agent networks—often within short, high-velocity windows that close before investigators can respond.

This is where one-time onboarding checks fall short. They confirm who opened an account, not who controls it months later. When identity is only verified at onboarding, identity farming remains one step ahead of detection.

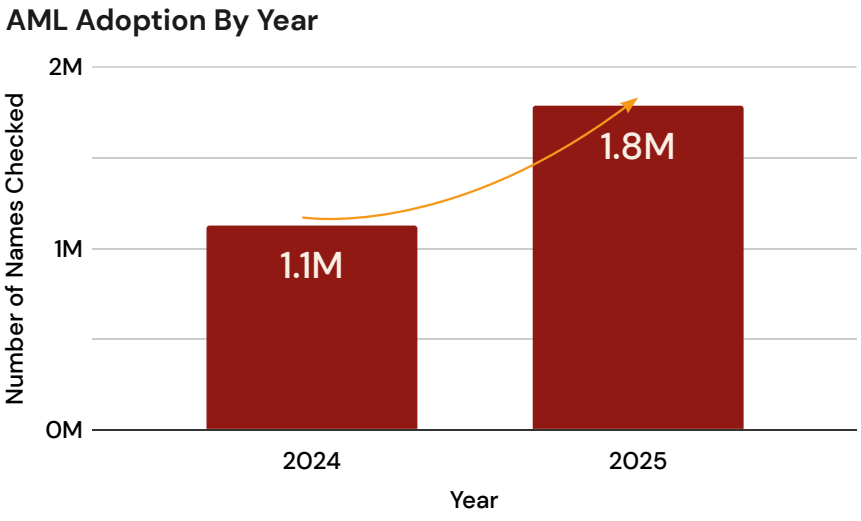

1.6 Fraud in Commerce

Not all losses come from synthetic identities. In 2025, a meaningful share of fraud in commerce came from chargebacks and promotion abuse, often driven by otherwise legitimate users exploiting dispute rules, issuer protections, and incentives. These patterns are typically seasonal or event-driven, which makes them more predictable—and more manageable—than novel technical attacks.

Chargeback & Seasonal/ Promotional Fraud

Loss often appears after verification because chargebacks are filed weeks later. This dynamic was visible in West Africa after the December 2024 travel/spending period: reporting in Nigeria and Ghana described a noticeable rise in chargebacks initiated from abroad during January 2025, aligned with "Detty December" / "Year of Return" travel.

Issuers often honor cardholder complaints, and smaller merchants may not have strong evidence linking the purchase to customer presence, so revenue earned during peak season can be reversed later through disputes.

Even without direct chargeback measurement, Smile ID signals show risk rising during peak commercial windows. During Black Friday in November 2025, duplicate ID checks increased by up to 500%, consistent with coordinated attempts to exploit promotions, bonuses, and account limits when traffic is high.

Daily Duplicate Checks - Betting Industry (Oct - Dec 2025)

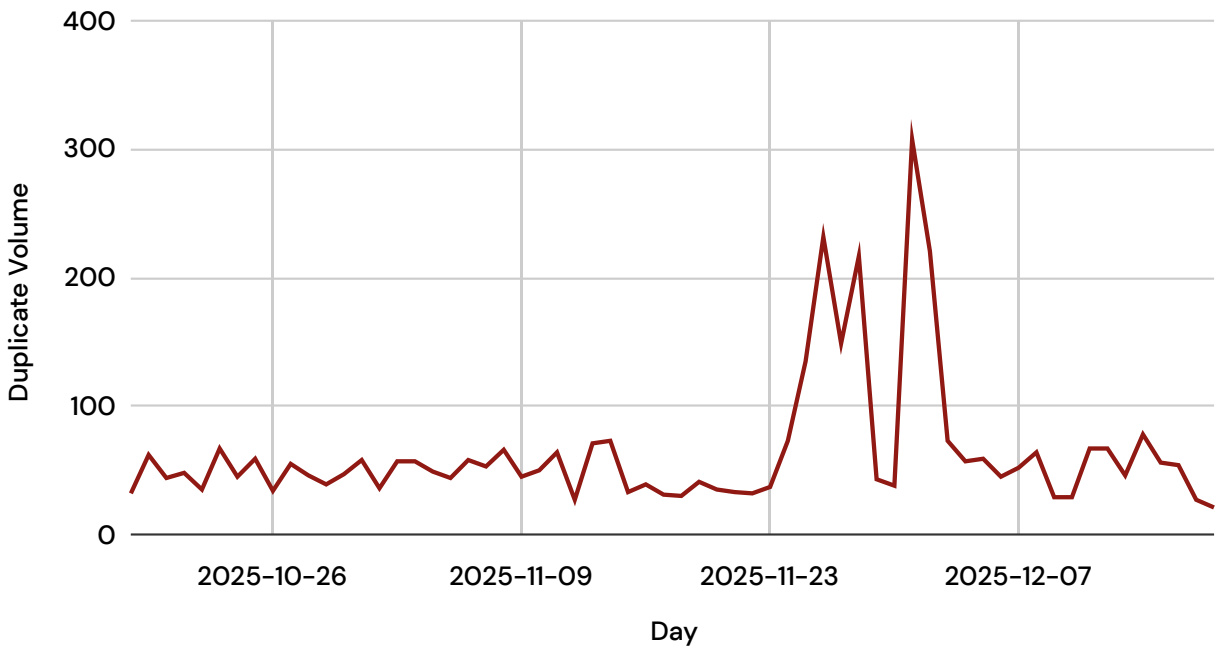

The operational risk is that volume creates blind spots for merchants and payment gateways: transaction review capacity gets stretched, responses slow, and repeat attackers gain time to reuse the same identity or card credentials across services.

In 2026, many preventable losses will come from fraud actors exploiting peak commercial periods when businesses are focused on growth and controls loosen under load.

Fraud in Context: Risks by Region

02

One face. Thousands of fraud attempts.

2.0 Africa is not a Single Risk Environment

The fraud techniques covered in Section 1—biometric manipulation, injection attacks, duplication, and identity farming—are showing up across many African markets. What differs is where these techniques are being used at scale, and which parts of the identity journey they target most. This section explains how the structural changes outlined in the introduction play out differently by region, shaped by identity infrastructure maturity, access to authoritative data, regulatory environments, and other factors.

While the percentage of overall verification attempts rejected for suspected fraud declined modestly between 2022 and 2025, the attempts became more targeted and high-impact, consistent with the trend toward sophistication and insider collusion. The relative percentage of fraud has declined, but with growing volume, the absolute number of fraud attempts continues to increase.

% of All Traffic Blocked for Suspected Fraud

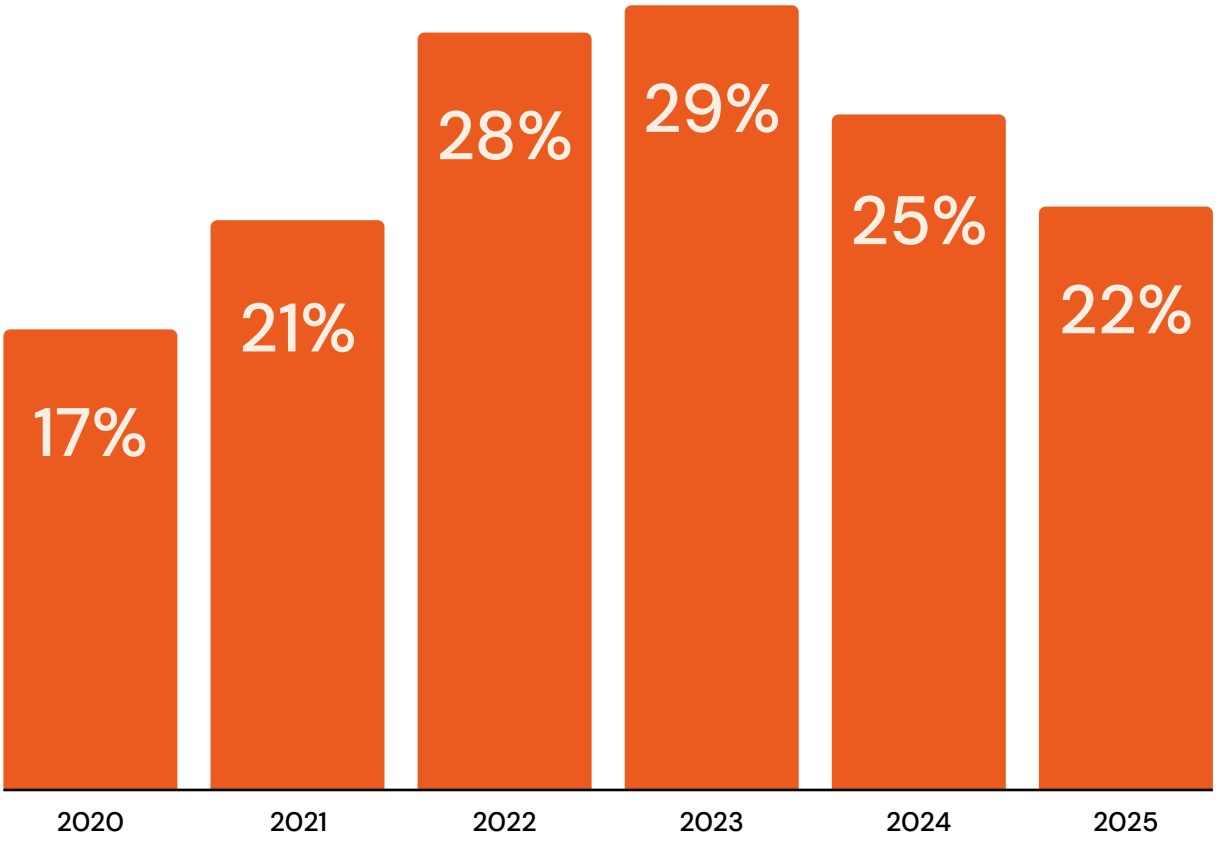

The relative % of fraud has declined, but with growing volume, the absolute number of fraud attempts continues to increase.

Regional Breakdown

- Central Africa
- East Africa
- Southern Africa
- West Africa

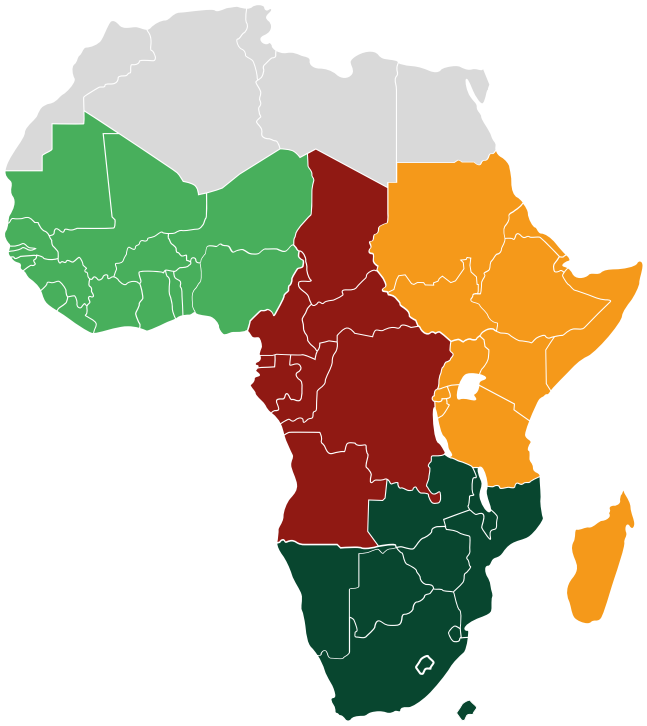

Fraud Rate By Region (2025)

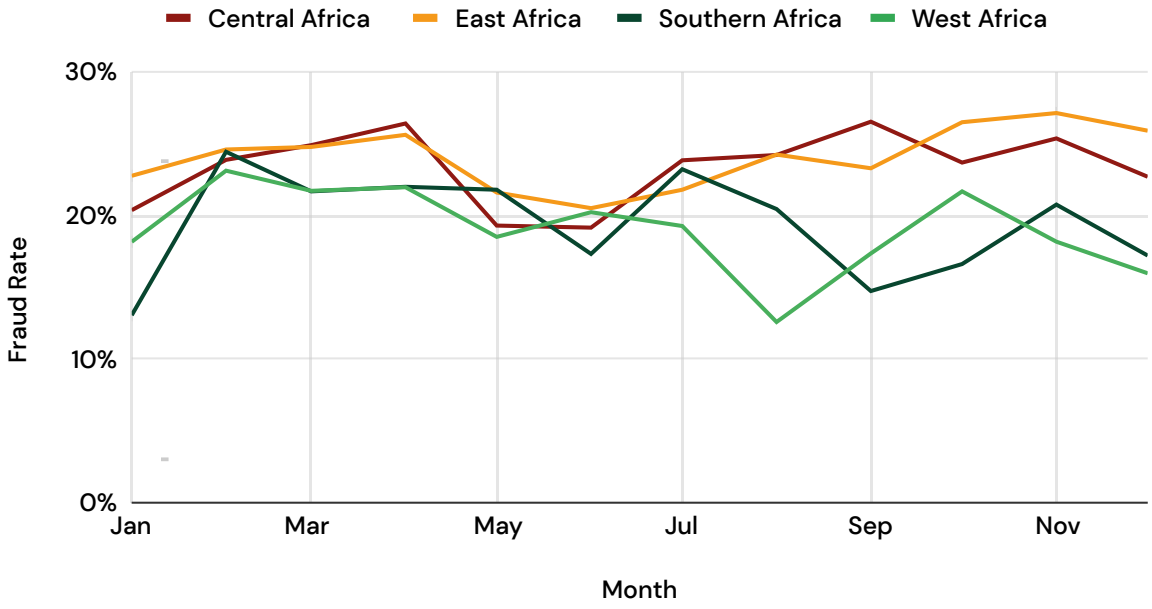

At a regional level, Central and East Africa are jointly highest overall, with 24% of verifications rejected for suspected fraud in each region. Both stay elevated through the year, typically fluctuating in the ~20–27% range.

In East Africa, most verification traffic comes from Kenya and Uganda, where national IDs are widely used but are not consistently enabled for fully remote, real-time biometric verification across all use cases. As a result, many onboarding and authentication flows remain document-led or hybrid, largely contributing to the high potential fraud rate in the region. A similar dynamic applies in Central Africa, where in several markets, verification remains document-heavy, and infrastructure for fully remote biometric verification is still maturing.

2.1 West Africa: Identity Farming & High-Value Account Takeover

Fraud in West Africa is increasingly driven by identity farming and account takeover, assisted by insiders. Attackers source or purchase identity assets, use mules or insider/agent-assisted onboarding to get accounts approved or take over existing high-value accounts, and then exploit those "trusted" accounts later—when they can move money quickly and cash out before controls react. The focus is less on creating new accounts and more on re-entering and abusing verified accounts at the moments where value concentrates.

Identity Reuse & Insider-Assisted ATO Dominates

Fraud rates in West Africa rise and fall throughout the year because attackers operate in campaign waves. Spikes tend to match periods when fraud groups launch coordinated runs, and dips often follow when they pause, rotate infrastructure, or change tactics.

West Africa Fraud Rate

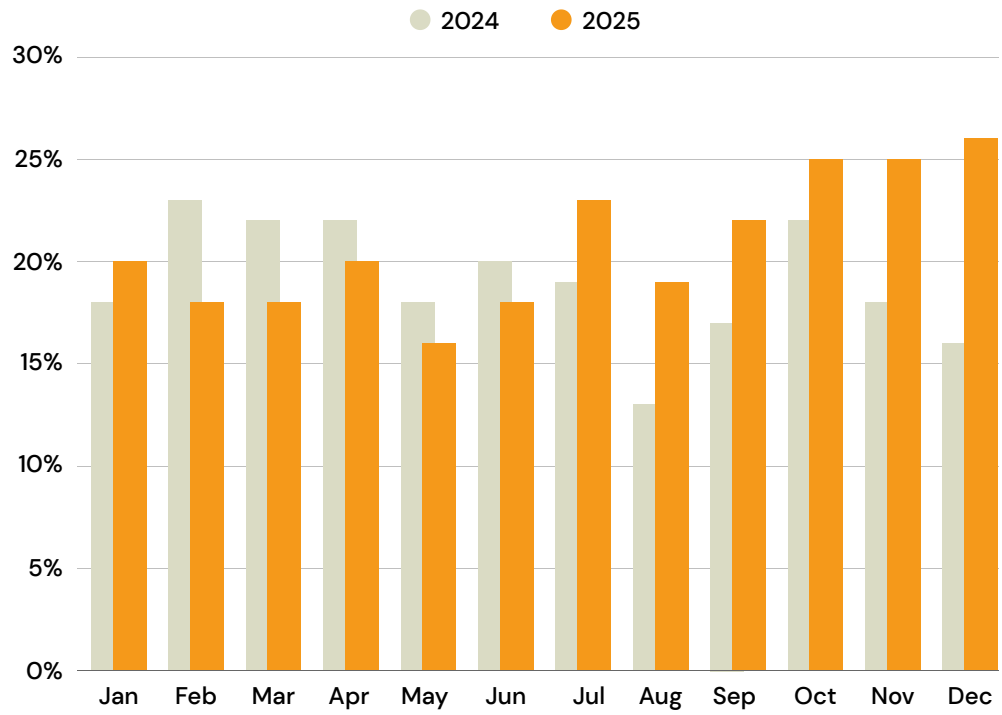

In 2025, regulators and enforcement continued to describe an identity supply chain: identity credentials and images are harvested or traded, then used to open, take over, or "weaponize" accounts at scale. The playbook starts with identity purchase and harvesting—real identity details and photos that can pass basic checks. These assets are used to create new accounts or to target existing accounts with higher limits. In Nigeria, warnings from the Economic and Financial Crimes Commission (EFCC) and ecosystem reporting pointed to large-scale harvesting of credentials such as BVN/NIN and passport photographs, reinforcing identity misuse as a key enabler of fraud.

Next is the enablement layer, often assisted by onboarding or insider help. Agents, intermediaries, or compromised staff help accounts get verified, reduce friction in KYC steps, or speed up access. This is where reuse scales: once a "working" identity package passes, it can be reused across multiple platforms.

After approval, attackers shift from onboarding to account takeover and re-entry, targeting the moments that unlock value: login and re-authentication, password/OTP resets, device changes, and transaction approvals. These steps are often designed for convenience, so they can offer an easier path into a trusted account.

Finally comes cash-out. With access secured, funds are moved quickly through high-velocity rails—previously created accounts, wallet transfers, crypto platforms, withdrawals, and cross-platform movement—often in short bursts so the money is gone before review or customer support can respond. FITC's Q1 2025 fraud reporting points to mobile/app-driven fraud as a leading loss channel, consistent with attackers prioritizing account access over older, noisier methods. Similar enforcement activity in Ghana highlights the same convergence of identity misuse, mobile-money fraud, and insider-assisted attacks.

Techniques that Enable Fraud in West Africa

Impersonation makes up about two-thirds (65%) of potential fraud attempts in West Africa, driven by spoofing and no-face-match during biometric verification.

Fraud Techniques in West Africa in 2025

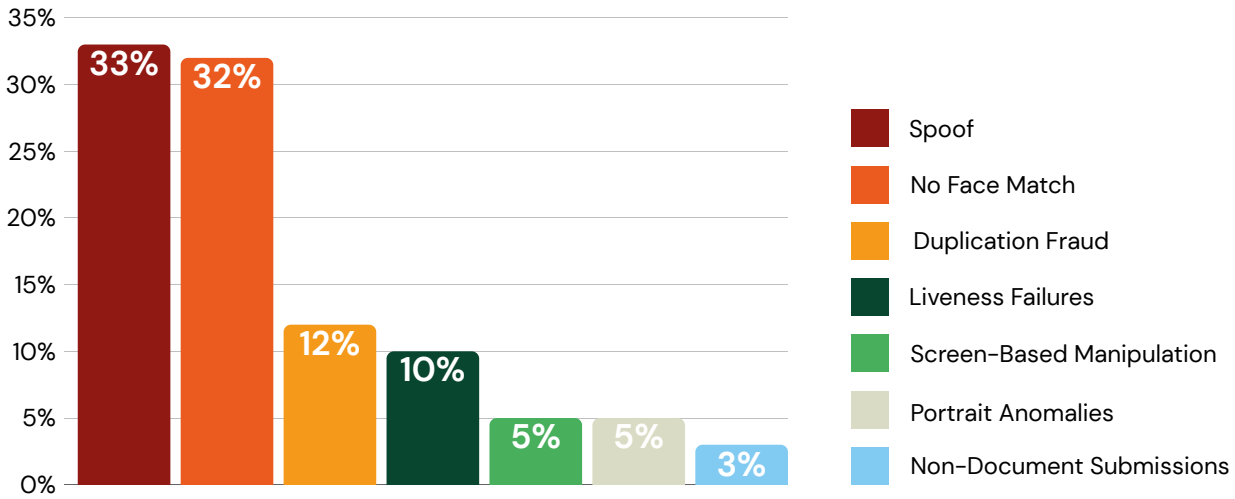

- **Impersonation (65%):** Spoof (33%) and no face match (32%) show attempts where the selfie is not tied to the claimed identity—linked to stolen identities, mules, and assisted or synthetic capture.
- **Reuse/automation at scale (22%):** Duplication (12%) shows the same biometric assets reused across sessions or accounts, often in bursts. Liveness failures (10%) point to replay, scripted attempts, or synthetic motion designed to bypass checks.
- **Document integrity & presentment tricks (13%):** Screen manipulation, portrait anomalies, and non-document submissions exist, but they are secondary.

This technique mix explains why biometric abuse is central in West Africa: a reused face or capture asset can be monetized repeatedly across platforms, making it more valuable than a one-time fake document.

Compliance Exposure Alongside Fraud Pressure

Alongside fraud, financial institutions in West Africa face a heavier operational burden from AML screening. In Smile ID's AML screening data, Ghana (~5%) and Nigeria (~5%) show the highest potential-match ratios—above South Africa (~4%) and far higher than Kenya (~1%) or Zambia (~1%).

Countries with Highest Ratio of Listed Persons

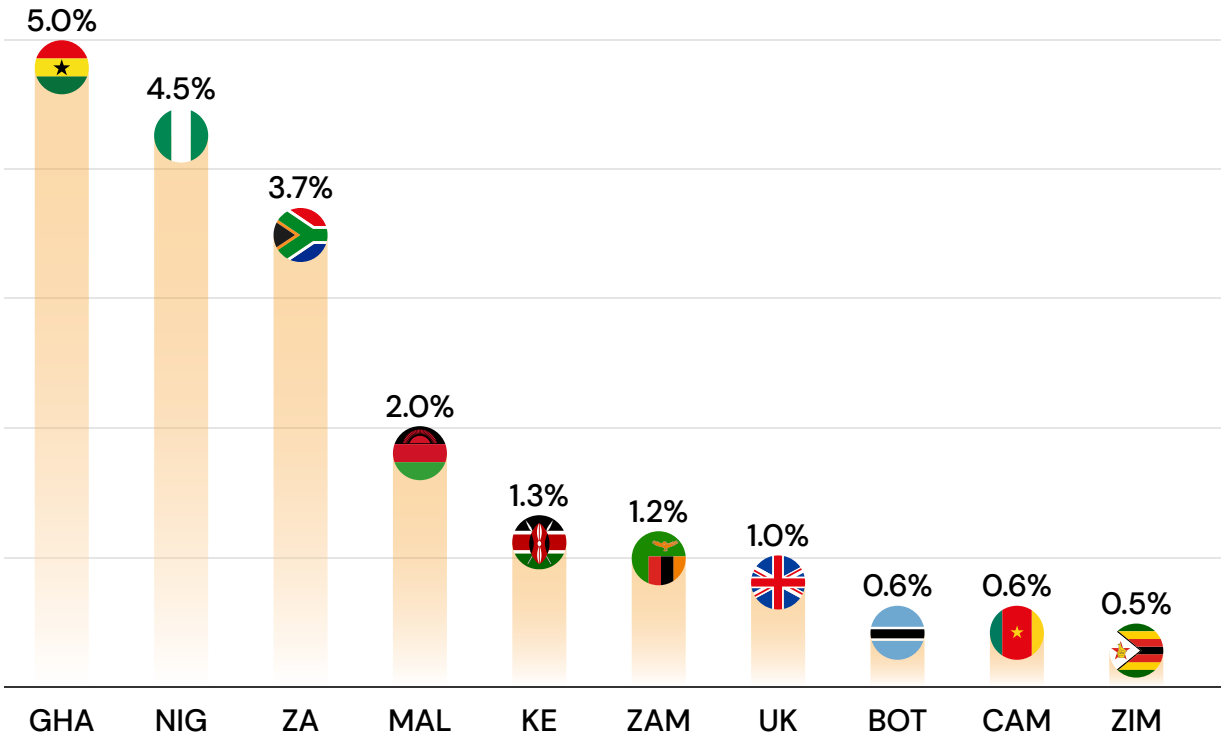

AML screening activity across Smile ID partner platforms. Data was normalized to account for relative market size and verification volume.

This does not mean West Africa has higher true sanctions or PEP exposure. It means more screenings return potential matches that require manual review and escalation, increasing workload, cost, and friction.

Fraud Follows Value: Shifts in Retail Banking

As mobile banking becomes more broadly adopted across Africa, fraudsters focus less on opening new accounts and more on targeting and taking over existing accounts. These attacks typically target verified users with transaction history and high stored value, account recovery paths (password resets, OTP flows), and login and re-authentication moments that are treated as "low risk."

In 2025, potential fraudulent attempts in retail-banking rose by 50% year over year. This increase is driven primarily by authentication and account-recovery events. Importantly, potential fraud attempts are growing faster than overall verification traffic in retail banking use cases—reinforcing that the observed change reflects attack behaviour, not just overall higher platform usage.

Retail Banking Fraud Rate in West Africa

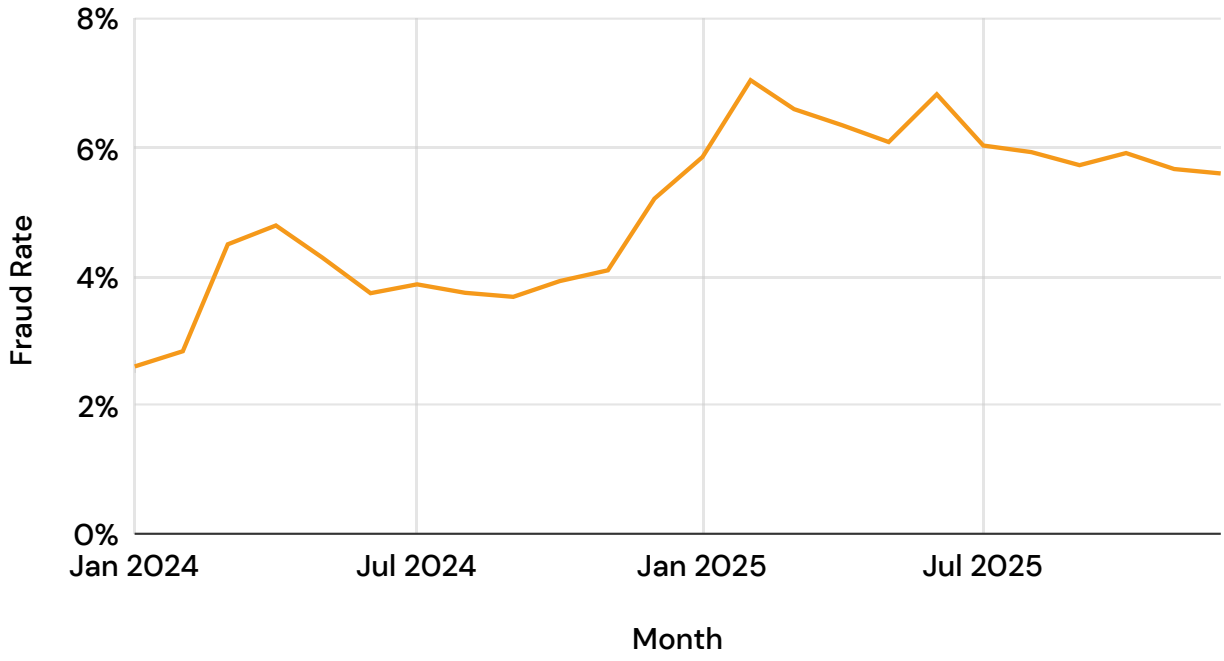

2.2 East Africa: Document Integrity in Hybrid Verification Journeys

In East Africa, many verification journeys are hybrid: users upload ID photos and sometimes verify against ID authorities, but ID database connectivity can be unreliable, and agents often help with onboarding. Because these flows rely heavily on documents, document fraud is a common risk.

In Kenya, Uganda, and Zambia, physical national IDs and user-submitted document images remain central to onboarding and account access. Even where national ID programs exist, remote, real-time biometric checks against authoritative sources are not consistent across services. Unlike in Nigeria or South Africa, there is not a "standard flow" for remote biometric onboarding. That creates a practical gap: most businesses rely on manual verification of documents or document-led flows where attackers don't need to beat sophisticated liveness algorithms if they can fool humans with forged documents—through portrait edits, photocopies, screen submissions, and borrowed IDs.

Why Documents Dominate in East Africa

In Kenya, a multi-agency crackdown in December 2025 targeted syndicates forging citizenship documents and IDs, with investigators warning that forged paperwork was creating financial exposure, including banks servicing loans tied to fraudulent documentation.

In Uganda, a 2025 attempted bank fraud case reported the use of a forged national ID to steal over $100,000. The common pattern is document-led impersonation: real or stolen IDs, edited images, or "clean" forgeries that require stronger image forensics and database cross-checks to detect.

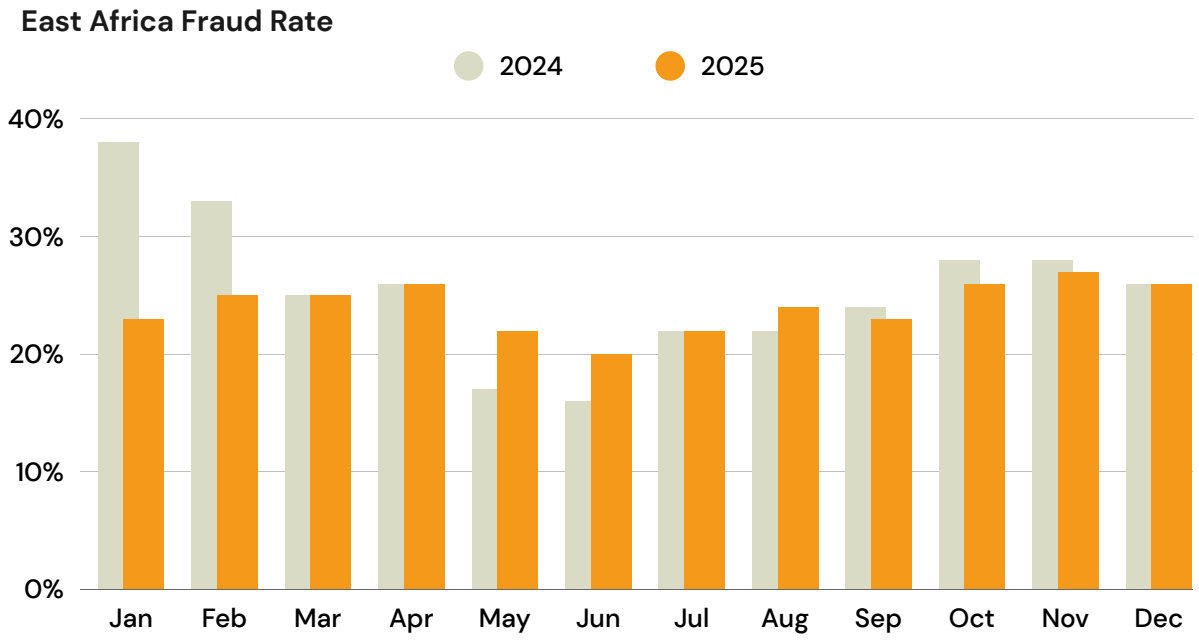

How Document Manipulation Facilitates Fraud in East Africa

In East Africa, about 3 in 5 verifications rejected for potential fraud are due to document integrity issues. The most common reason for rejection is portrait anomalies—face swaps or insertions that preserve the document layout. These represent roughly a third of all rejections in the region.

In East Africa, fraud surfaces when verification systems attempt to bind those documents to a real, present human.

Fraud Techniques in East Africa in 2025

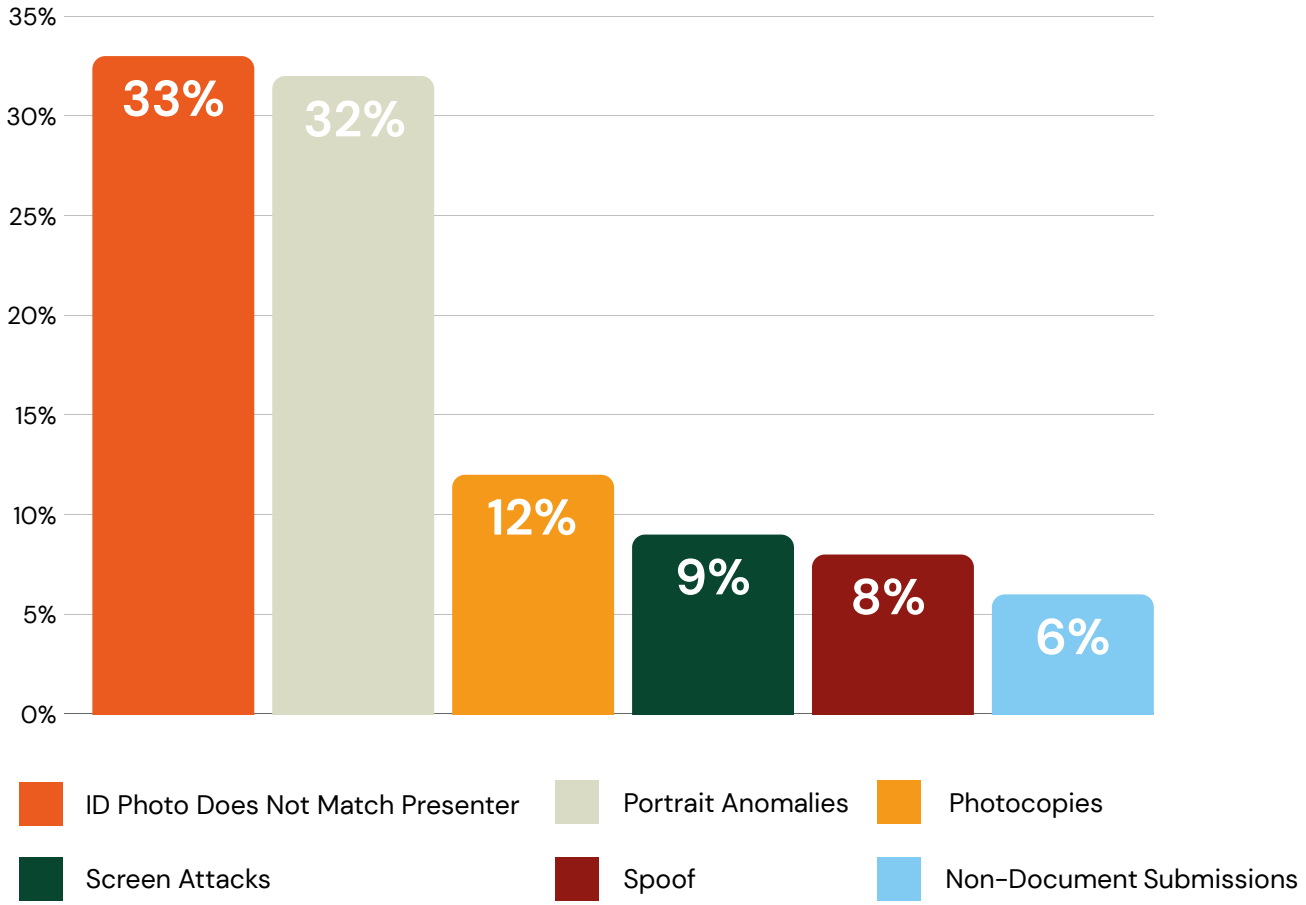

- **Document integrity & presentment (60% of rejections):** Portrait anomalies (33%), photocopies (12%), screen attacks (9%), and non-document submissions (6%)
- **Impersonation (33%):** The live selfie does not credibly match the claimed identity on the document—often signs of borrowed IDs, mule-assisted onboarding, or portrait manipulation
- **Biometric spoofing (8%):** Attempts to defeat liveness checks. So far these attacks have been less common in East Africa than in West, but they are growing in frequency and sophistication

Where Fraud Concentrates

Based on 2025 averages in Smile ID's East Africa traffic, the highest fraud risk segments are cryptocurrency (37%), payments (35%), and investment (32%). Cryptocurrency remains consistently high, with pronounced spikes in early-year and late-year windows, aligning with increasing regulatory focus on virtual asset providers.

% of Verifications Rejected for Suspected Fraud by Industry (East Africa) in 2025

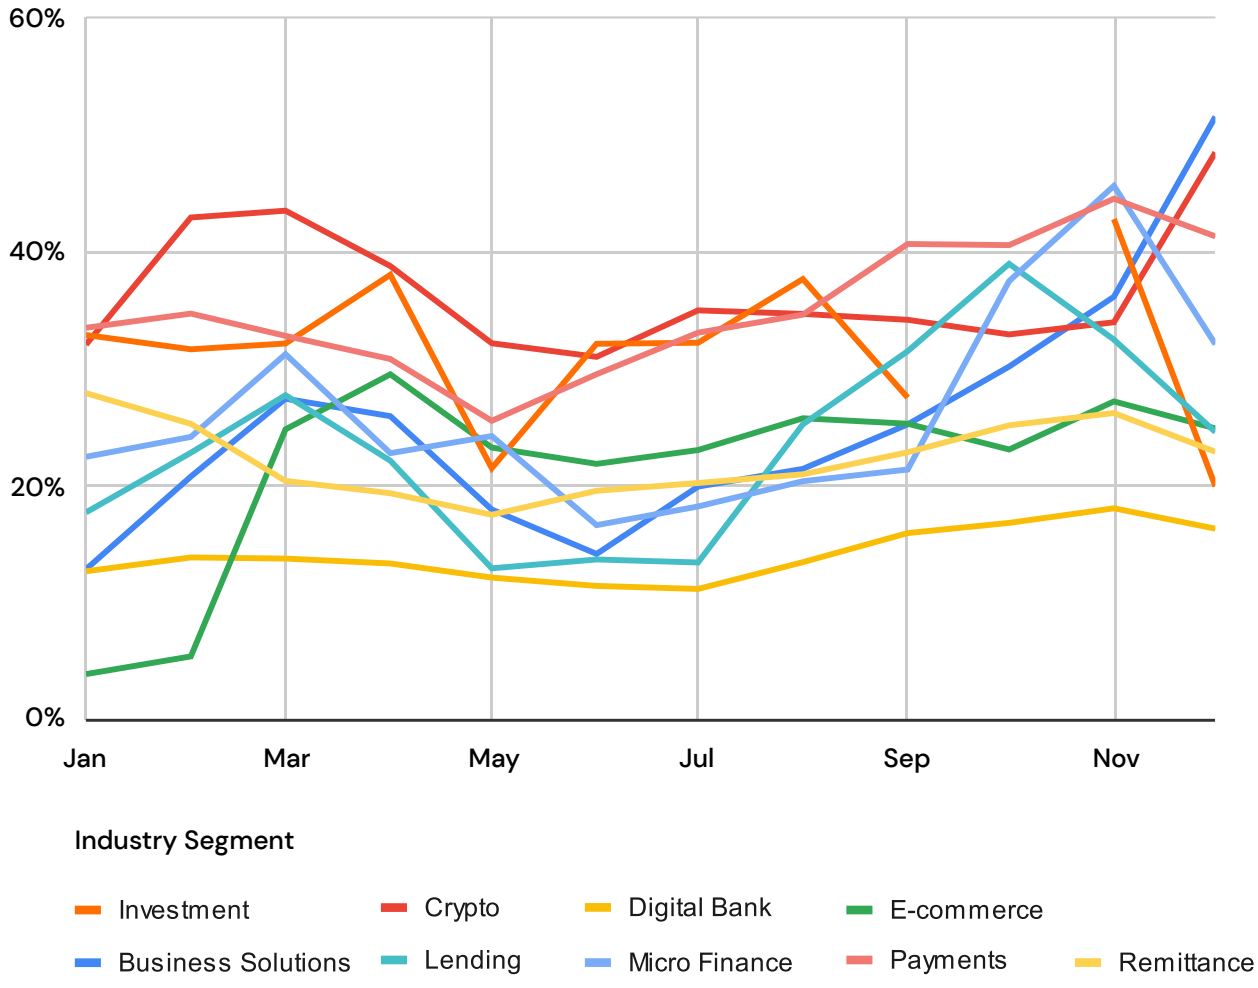

How Infrastructure Shapes Fraud Concentration

In parts of East Africa, platforms must adapt to real-world constraints such as intermittent connectivity, shared devices, and agent-assisted journeys. These improve inclusion, but they also widen the gap between capture and validation, especially when checks are delayed or asynchronous.

Channel mix matters too, e.g many users still rely on USSD for financial services (e.g., Safaricom supports M-PESA via USSD), or use WhatsApp as a primary communication channel, which can limit consistent enforcement of biometrics and secure capture, increasing reliance on documents for verification.

2.3 Southern Africa: AI-Assisted Impersonation and Spoofing

Fraud attacks in Southern Africa are primarily focused on biometric manipulation techniques. Like in most mature digital markets, selfies are not just for onboarding—they are used repeatedly for login, recovery, and high-risk account actions. That makes biometrics the shortest path to account access, and it is where attackers concentrate.

Southern Africa Fraud Rate

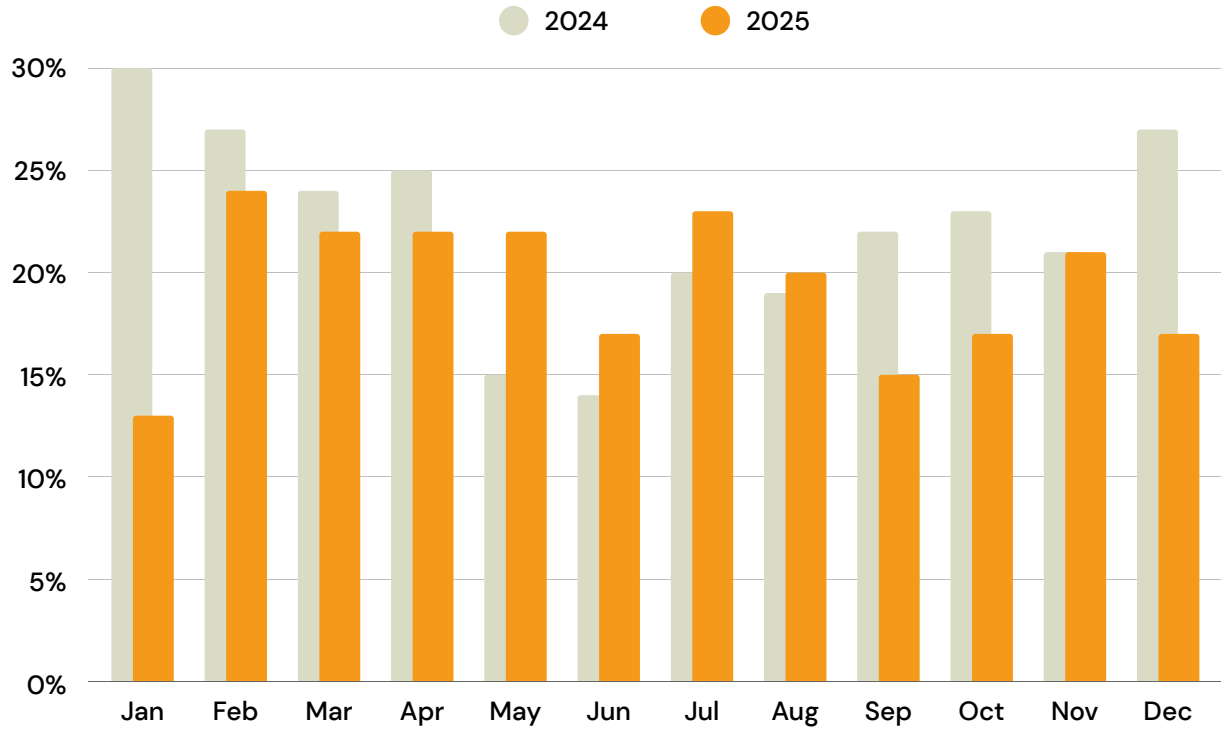

AI-Assisted Impersonation as the Dominant Threat

Across Southern Africa in 2025, regulators and central banks repeatedly warned the public about AI-generated impersonation used to push fraudulent financial offers. In South Africa, the South African Reserve Bank warned about a deepfake video impersonating Governor Lesetja Kganyago to promote a fake investment opportunity. Zimbabwe's Reserve Bank also issued a public alert about an AI-generated scam promoting a fake investment platform that misused the governor's identity.

Zambia's central bank similarly warned about impersonation of its governor in the promotion of a fraudulent financial scheme. South Africa's FSCA also issued public warnings about AI-linked impersonation and deceptive investment promotions during 2025.

Why Biometric Impersonation Dominates Fraud in Southern Africa

Fraud is overwhelmingly biometric in South Africa: nearly 9 in 10 verification attempts rejected for potential fraud are due to impersonation and spoofing (47% no-face-match and 40% spoofing) during biometric verification.

Fraud Techniques in Southern Africa in 2025

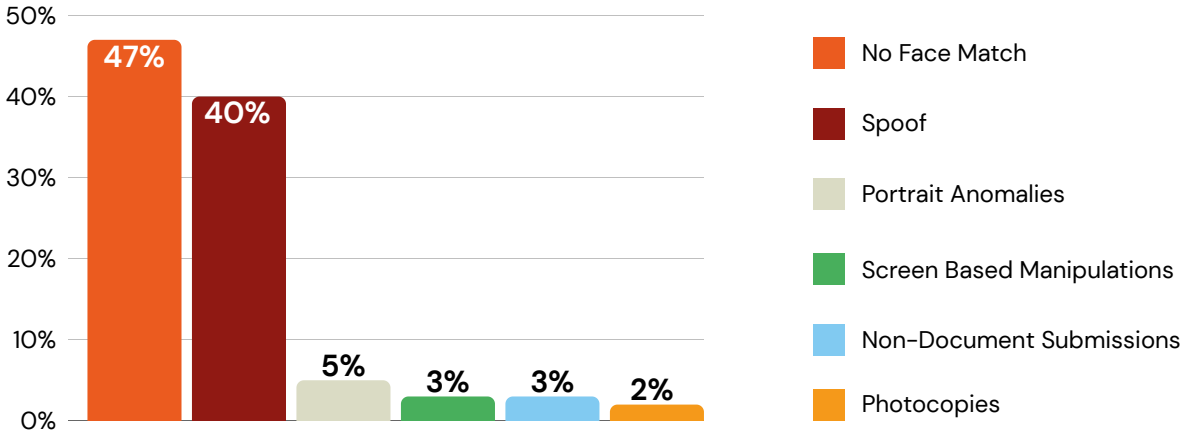

- **Impersonation (47%):** No-face-match dominates, meaning the person presenting cannot be credibly linked to the claimed identity—consistent with stolen identities, mule-assisted flows, and post-onboarding abuse.
- **Biometric integrity attacks (40%):** Spoofing attempts aimed at defeating liveness and similarity checks, including deepfake-assisted impersonation and face-swap techniques.
- **Document integrity & presentment tricks (13%):** Present, but clearly secondary.

This is reinforced by the rate of change: deepfake attempts averaged 2% in 2025 but surged from less than 200 monthly attempts in 2024 to over 3,000 attempts per month by year end.

Rate of Spoof attempts in Southern Africa

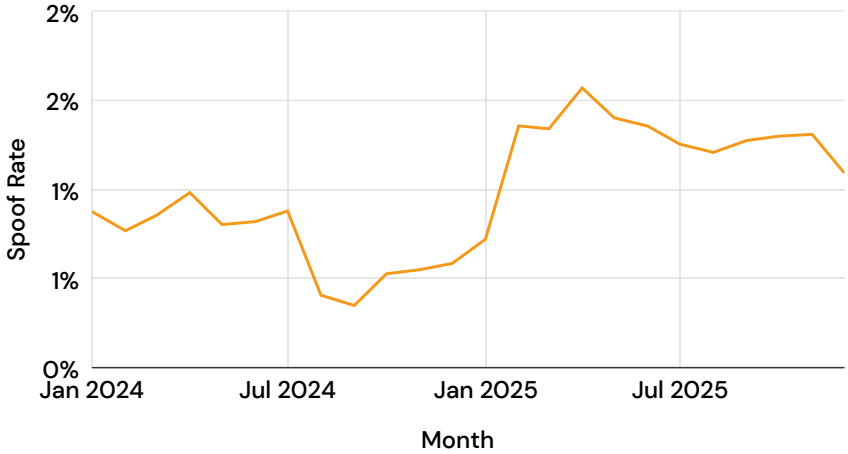

Chart shows confirmed spoof rate across Smile ID partner platforms in Southern African region between January and December 2025.

BNPL & Digital Lending: Speed and Month-End Fraud Pressure

BNPL and digital lending attract fraud because approvals are fast and low-friction. With limited time for strong upfront checks—and a gap between approval, purchase, and repayment—attackers can use stolen or reused identities to open or take over accounts, then reuse the same faces, devices, and workflows across multiple lenders before controls catch up. Common patterns include biometric spoofing in fast onboarding flows, first-loan abuse using previously used or compromised identities, and identity reuse across multiple lenders.

Fraud attempts tend to peak in week four of each month as personal budgets tighten.

BNPL Fraud Rate by Week in 2025

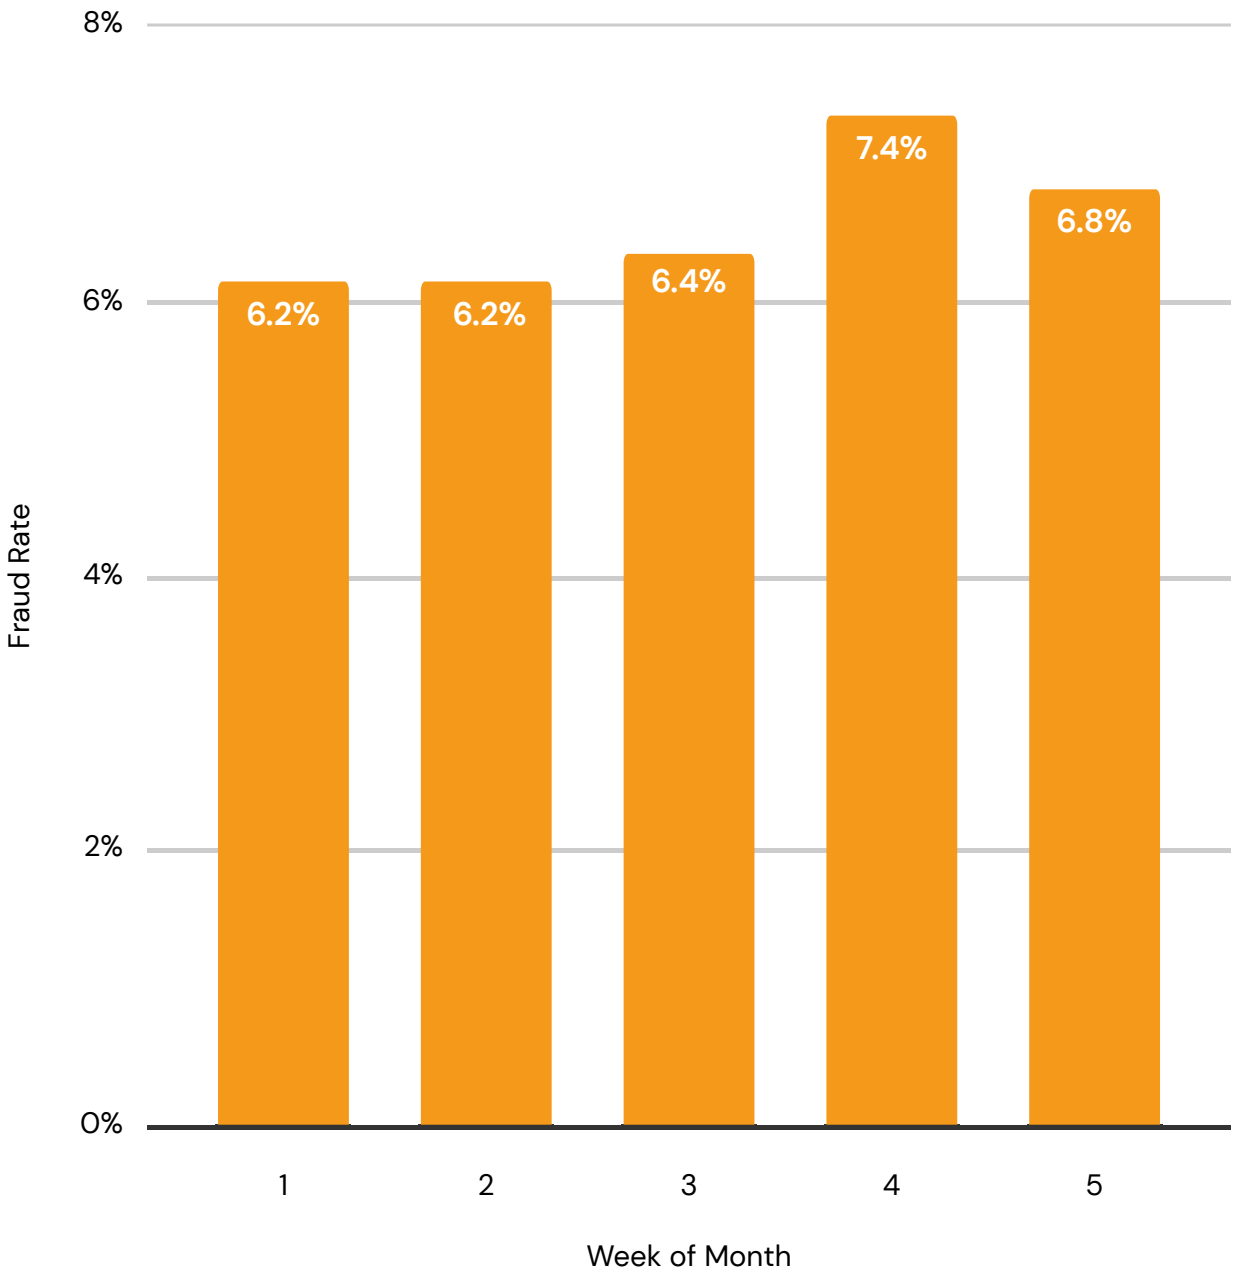

Gaming, Betting, and Digital Commerce: Promotion-Linked, Repeatable Abuse

These platforms attract fraud when incentives reward sign up. Bonuses, referrals, and promotions are harvested repeatedly using the same identity assets.

Daily Duplicate Checks - Betting Industry (Oct - Dec 2025)

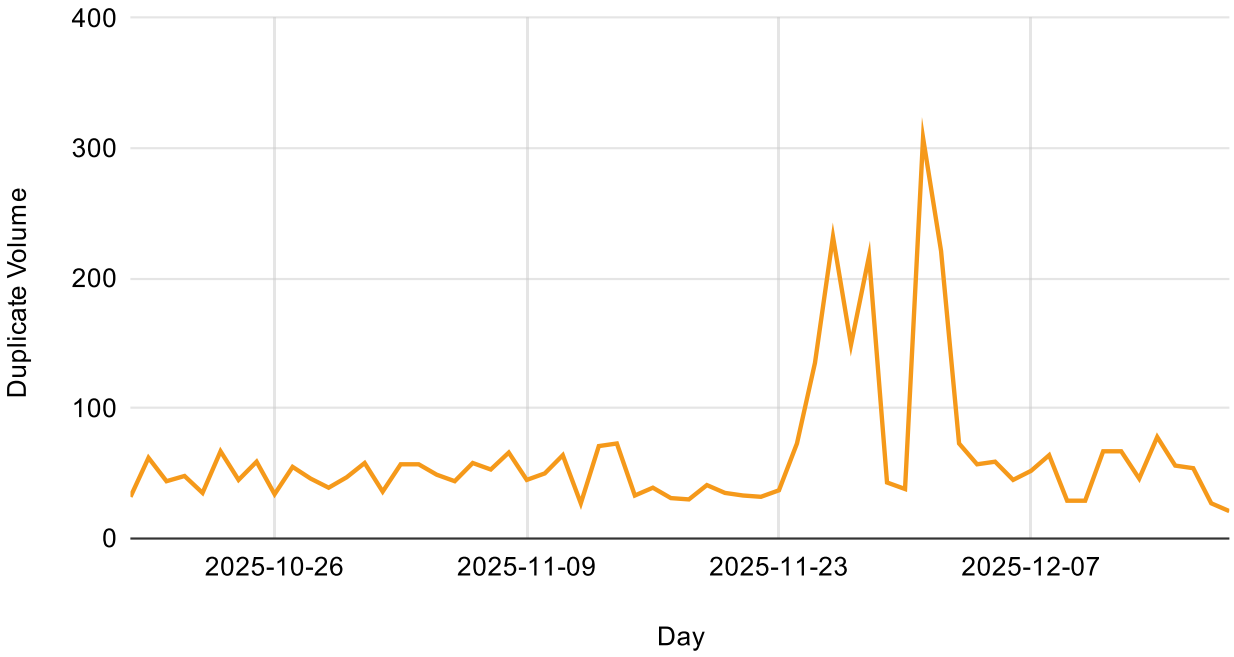

Smile ID data shows how quickly this abuse can scale. During Black Friday 2025, duplicate identity checks in betting platforms increased by 500%, rising from an average of 50 daily attempts to peaks above 300.

Common patterns include:

- Bonus/referral fraud through coordinated account creation.
- Eligibility resets via duplicate accounts and identity recycling.
- Aggressive reuse of faces, devices, and environments (including emulators).

Where Value Concentrates in Southern Africa

Based on Smile ID's Southern Africa traffic, the highest-risk segments are remittance (20%), digital banks (15%), and crypto (13%).

% of Verifications Rejected for Suspected Fraud by Industry (Southern Africa) in 2025

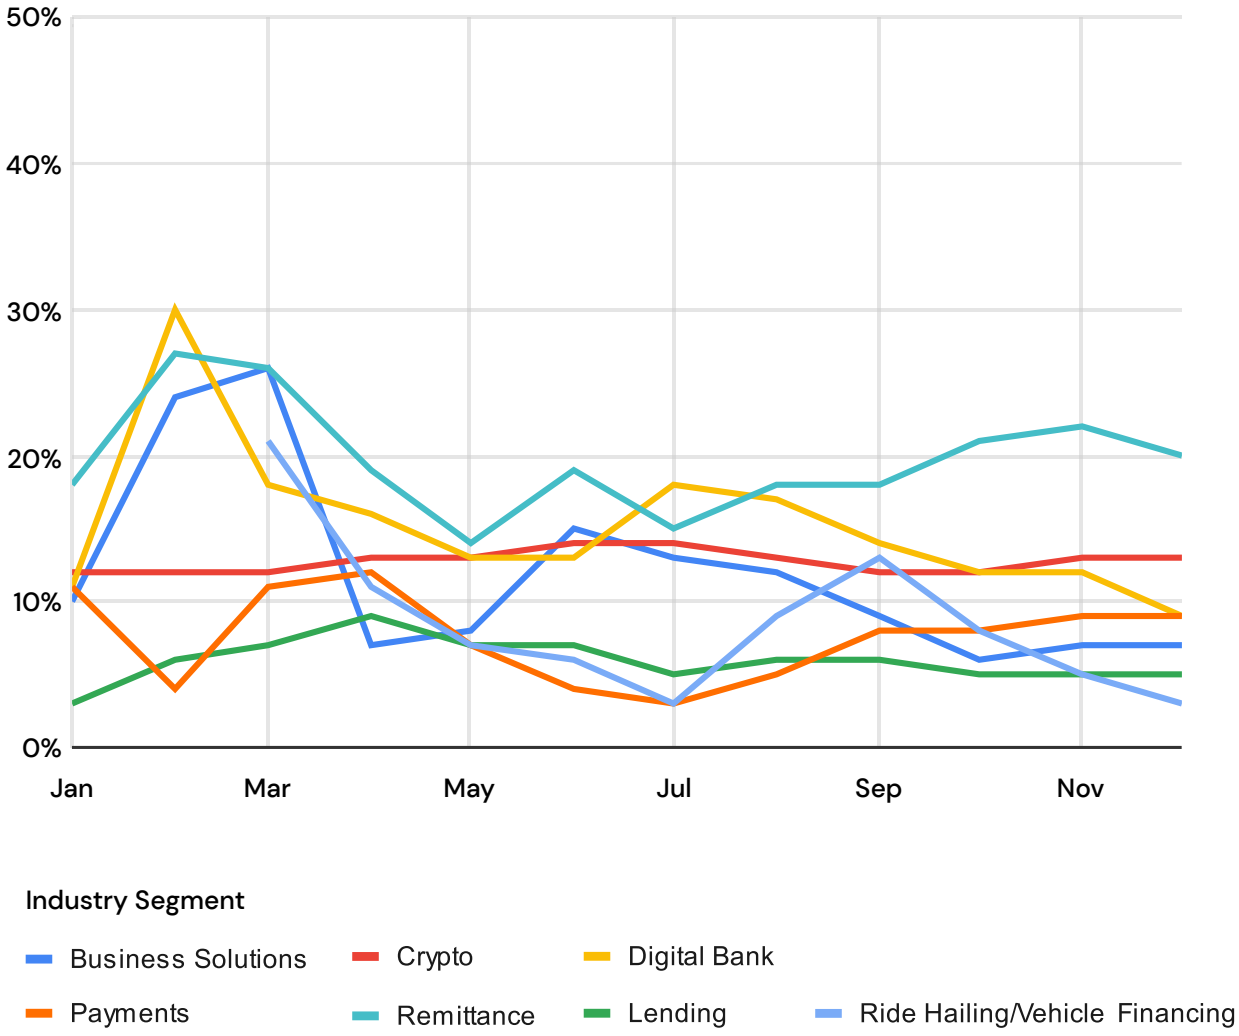

2.4 Francophone Africa: Document Reliance Under Regulatory Constraints

While francophone countries differ significantly in market operation, infrastructure, and industry mix, many share structural characteristics that shape how identity verification operates in practice. In several markets, regulatory frameworks place restrictions on biometric data processing, which can limit private-sector access to authoritative biometric matching and increase reliance on document-based verification flows.

Impersonation & Identity Reuse as Primary Threats

Where verification flows rely primarily on document inspection and biodata matching—without consistent ability to confirm biometrically that the person presenting a document is its rightful holder—attackers don't need advanced techniques to succeed. Credible-looking documents combined with identity reuse are often sufficient to enable impersonation and abuse.

Public reporting shows that document-falsification networks remain active in these markets. Cases reported in 2025 involved the manufacturing of false administrative and identity documents in Senegal, and heightened government concern and coordination against identity and document fraud in Côte d'Ivoire. Fraudsters don't need "perfect deepfakes" to win—they can win with credible documents and weak proof-of-ownership controls.

Techniques Driving Fraud

In Francophone markets, fraud is mostly document-driven. About two-thirds of all rejected verifications were due to document fraud—led by portrait anomalies (22%) and photocopies (20%)—with no face match (26%) showing impersonation when the person presenting the ID is not the rightful owner.

Fraud Techniques in Francophone Africa in 2025

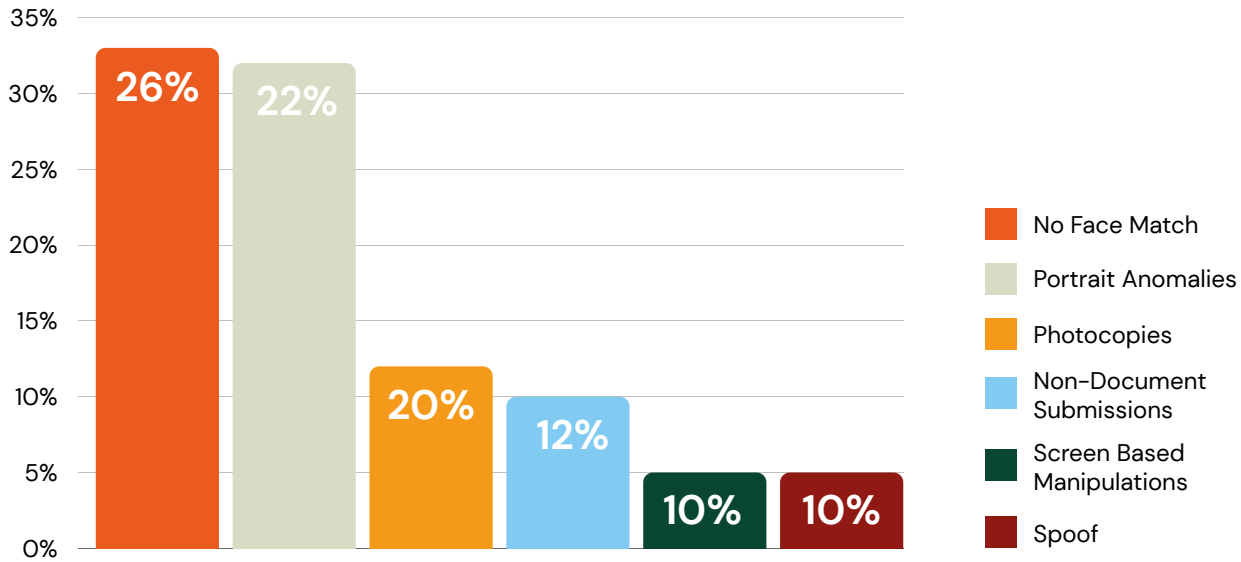

Chart shows the distribution of fraud techniques across Smile ID verifications in Francophone markets. It illustrates relative attack patterns under authority-data constraints, not uniform behaviour across countries.

- **Document integrity & presentment tricks (64%):** Portrait anomalies and photocopies reflect altered or reused document presentation, while screen-based manipulation and probing behaviour suggest testing for workflows that will accept "good-enough" submissions
- **Impersonation (26%):** No-face-match signals indicate attempts where the selfie input does not credibly link to the claimed identity, consistent with impersonation rather than biometric spoofing
- **Biometric integrity attacks (10%):** Spoof signals are present but secondary, reinforcing that observed risk is driven more by document-based impersonation than by advanced biometric attack techniques

This is another form of the same systems problem described earlier in this report: when strong authority signals aren't available, defence depends on connecting weaker signals over time to expose reuse and coordination.

2.5 Central Africa: Provenance Gaps, Assisted Capture & Document Exploitation

Observations in this section are primarily from verification traffic in Cameroon, Gabon and the DRC, where assisted onboarding, retries, and delayed submission remain common parts of identity workflows. Because many IDs and selfies are captured outside a tightly controlled app flow, there is often no strong proof of where, when, or how the images were captured—even if they look real.

When there is a long delay between capture and validation, it becomes easier to swap what gets submitted: a photocopy instead of an original, a screen replay instead of a live document, or an assisted impersonation using someone else's ID. As a result, fraud in these markets focuses more on document reuse and assisted impersonation than on beating the biometric model itself.

Why Provenance Gaps Create Fraud Risk

Where connectivity is uneven, platforms rely on assisted capture and deferred uploads. The longer the delay, the larger the opportunity to substitute or replay evidence. In the DRC, internet access remains limited (around 30% online), and the government's 2025 move to license satellite internet reflects continued connectivity gaps.

In Gabon, ARCEP has reported persistent coverage deficits across roughly 1,253 villages, reinforcing why assisted onboarding and deferred submissions remain common outside major urban centers.

In Central Africa, no face match during document verification accounts for over a third of all fraud rejections—the single largest category. Document presentment attacks (photocopies, screen-based manipulation) and portrait anomalies together represent the majority of observed fraud, while biometric spoofing remains comparatively low.

Automated document forensics and capture validation can easily strengthen controls in these workflows improving speed and customer satisfaction without adding friction, particularly where manual review capacity is limited.

Techniques Driving Fraud in Central Africa

In Central Africa, the single biggest potential fraud rejection reason is no face match during document verification (36%)—about 1 in 3 flagged cases.

- **Document presentment attacks (37%):** Non-document submissions, photocopies, and screen-based manipulation together reflect substitution and replay attacks.
- **Impersonation (36%):** No-face-match rejection reason accounts for the single largest category, indicating attempts where the selfie input does not credibly link to the claimed identity — often consistent with assisted presentation or reused documents rather than biometric spoofing.
- **Document integrity attacks (22%):** Portrait anomalies indicate targeted portrait-region tampering that preserves the rest of the document structure, allowing altered documents to pass superficial checks.
- **Biometric integrity attacks (4%):** Spoof signals remain comparatively low, reinforcing that observed risk in Cameroon and Gabon is driven primarily by document and presentment abuse rather than advanced biometric impersonation.

Fraud Techniques in Central Africa in 2025

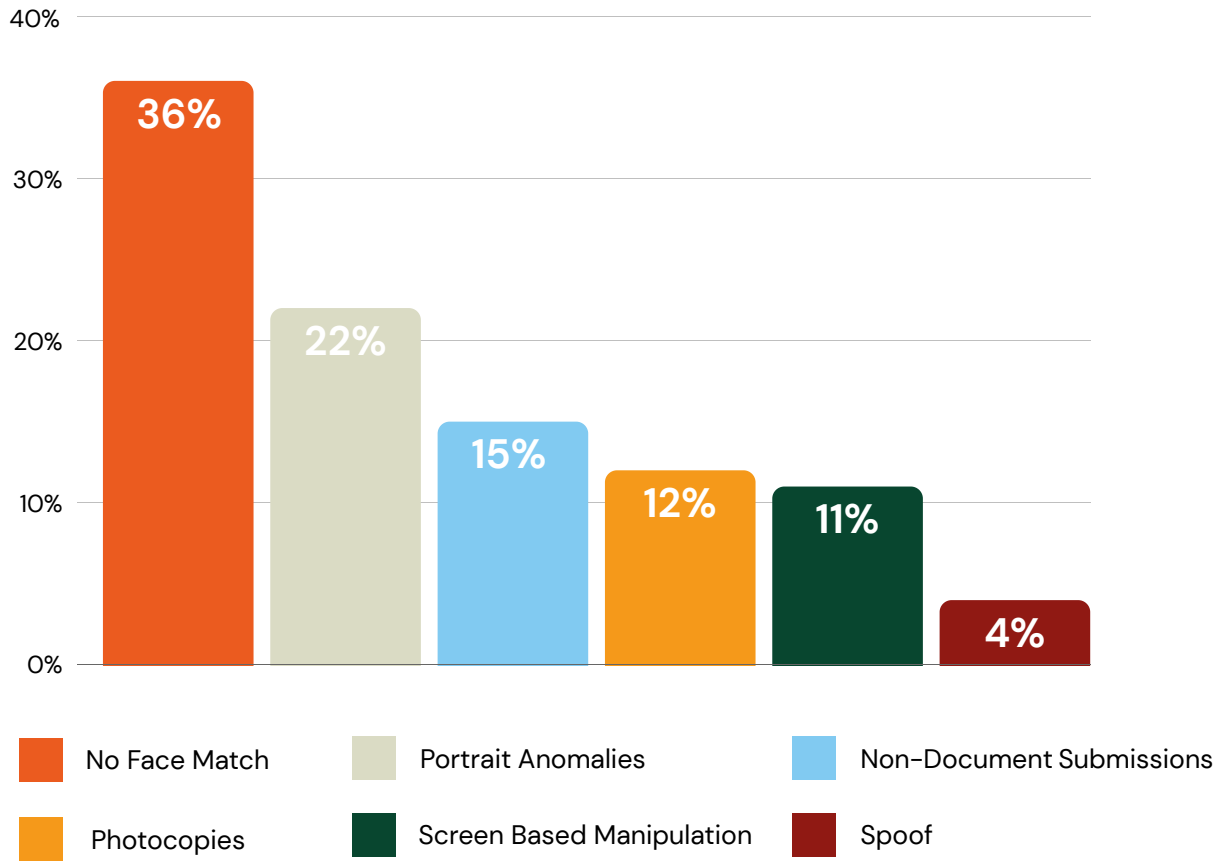

The Foundation: How Modern Verification Works

03

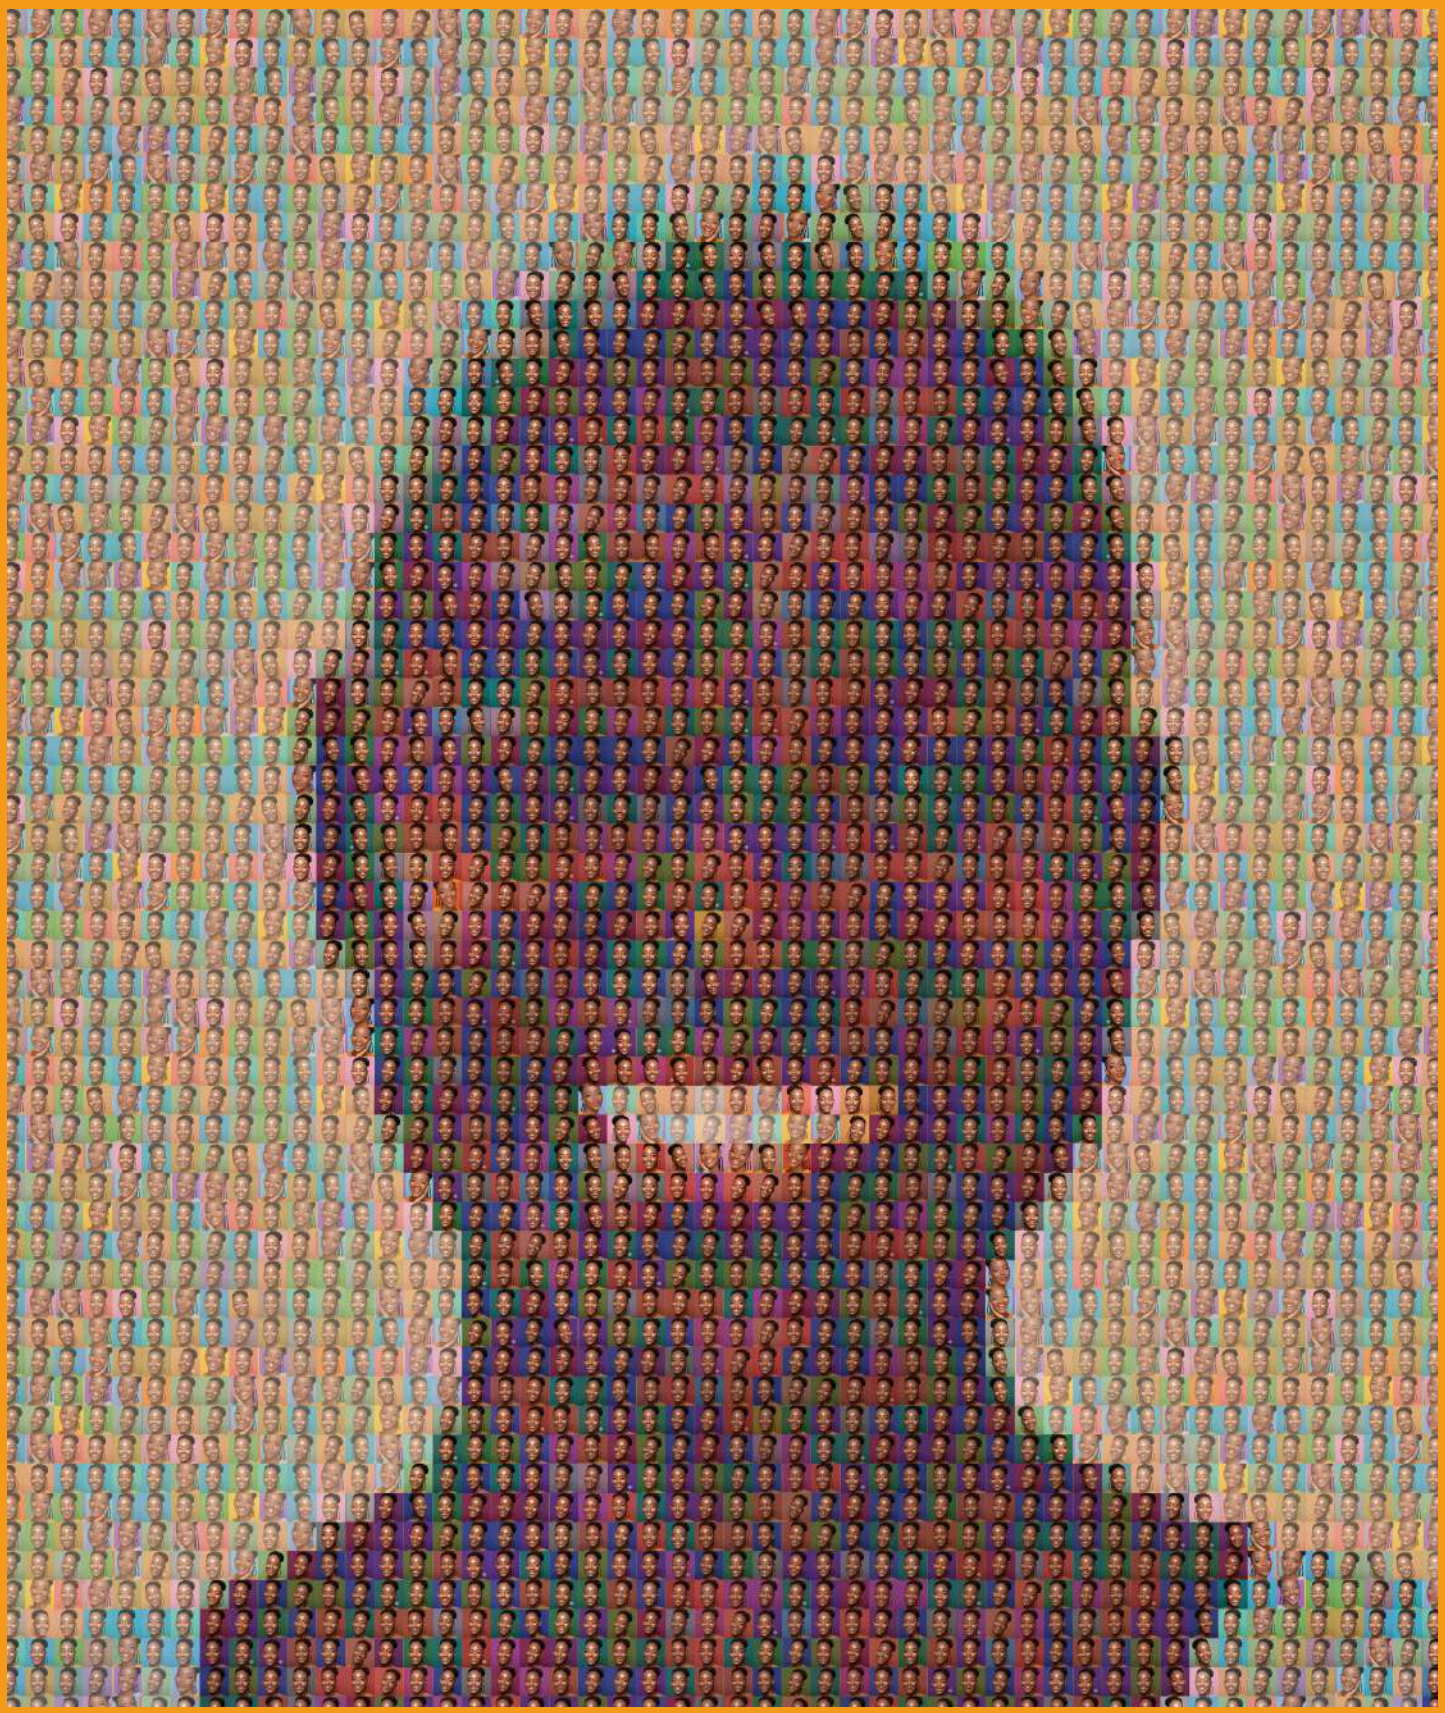

One face. Thousands of fraud attempts.

Effective fraud defence begins with reliable verification—the ability to determine whether an attempt is live, technically plausible, and linked to a real identity. But as fraud becomes more sophisticated, verification systems face a fundamental tension: they must catch evolving attacks without rejecting legitimate users. Single-method approaches struggle with this balance.

Why Smile ID Combines Classic Detection and AI

Legacy algorithms consistently catch known, repeatable fraud patterns—spoofs, replay attacks, and document manipulation techniques that have been seen before. But they can miss novel variants. Multi-modal large language models, by contrast, surface emerging patterns and adapt to novel attacks. But they can also increase friction by falsely rejecting legitimate transactions—flagging edge cases or unusual-but-legitimate capture conditions as suspicious. Smile ID combines both approaches: legacy controls remove repeatable abuse reliably, while adaptive AI keeps pace with novel attacks. This layered strategy improves detection while avoiding unnecessary friction for genuine users. When a new attack technique emerges, AI flags it. Once confirmed as fraud, it feeds into the legacy detection layer so future attempts are blocked automatically without requiring AI inference—reducing latency and false positives while maintaining strong security.

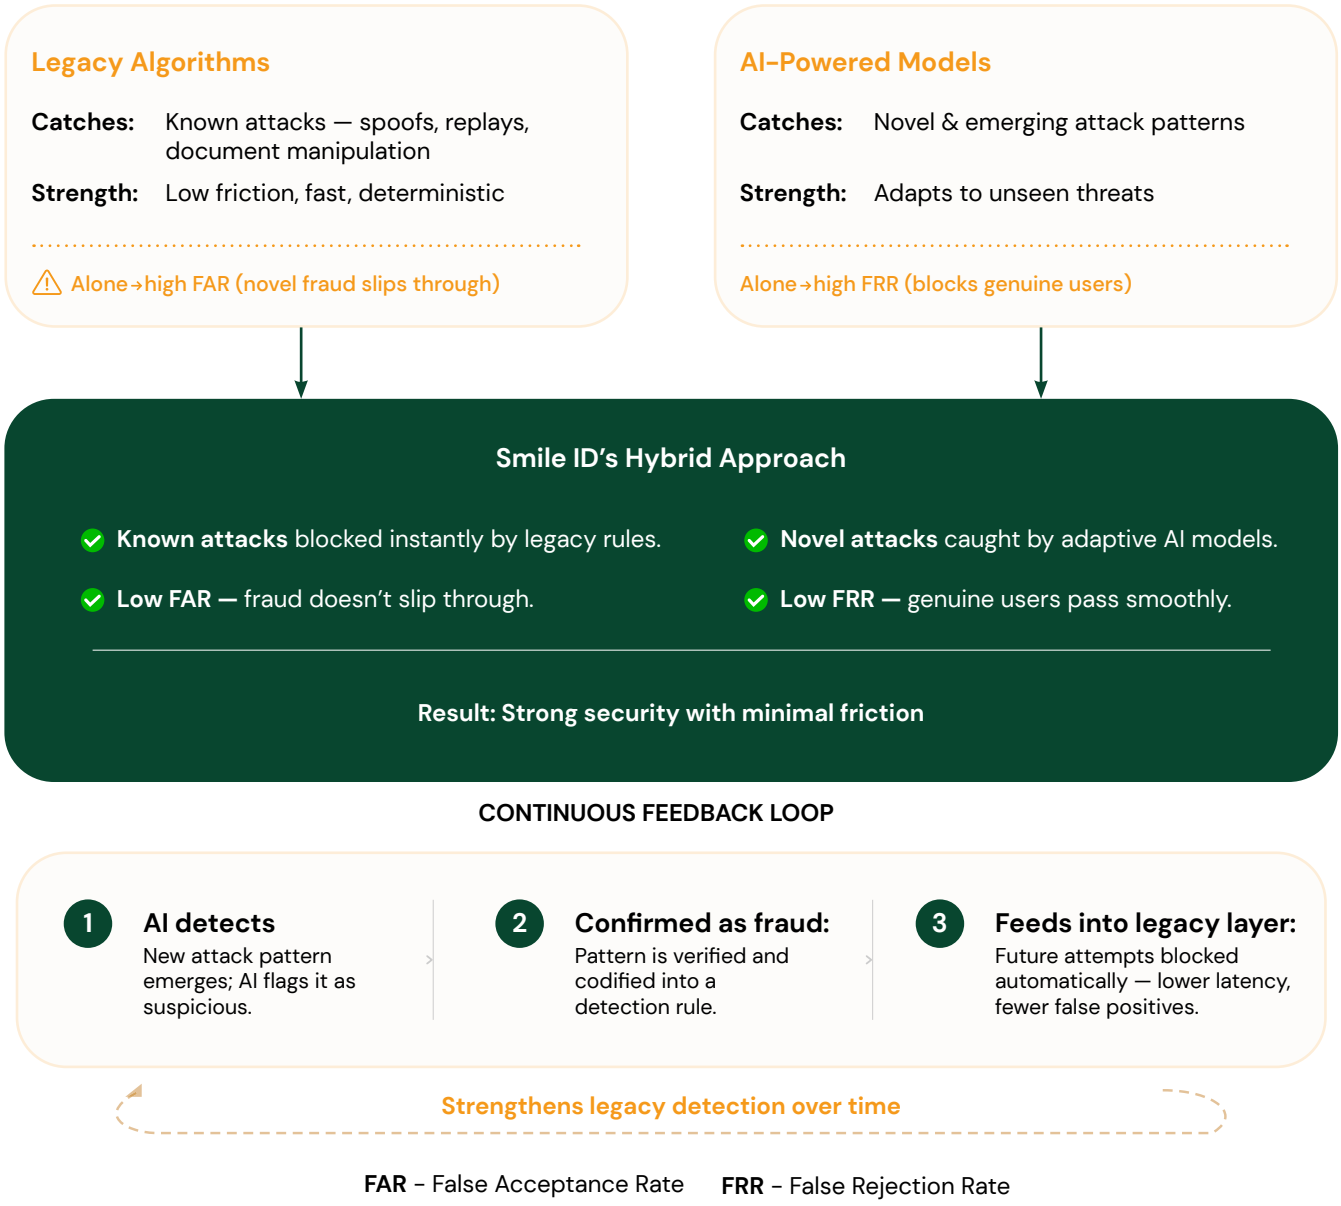

FAR - False Acceptance Rate **FRR** - False Rejection Rate

How Verification Layers Work Together

Smile ID's verification system operates across multiple detection layers, each designed to catch specific fraud types. The distribution below shows where suspected fraud is caught across these layers in 2025:

Smile ID Fraud Check Models

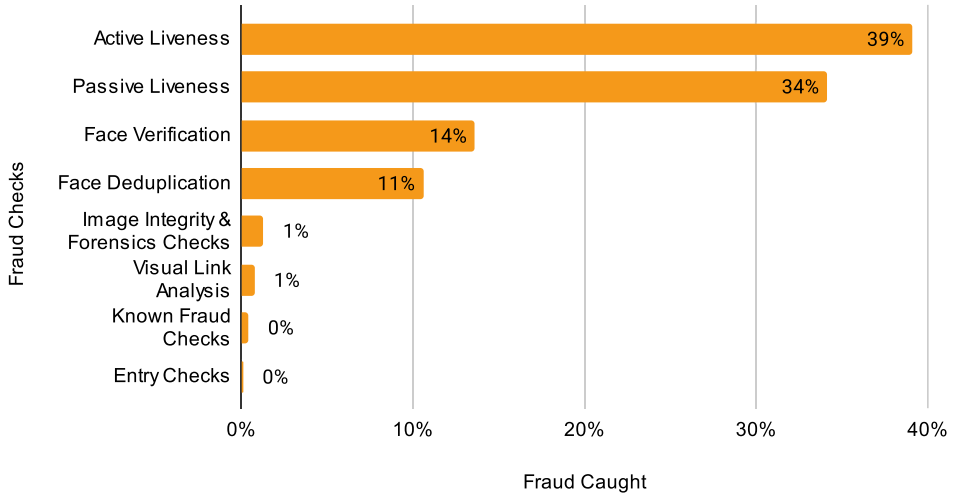

Entry Controls: Blocking Known Abuse (<1%)

Block known bad assets and environments before deeper analysis begins—basic hygiene that stops repeat offenders immediately.

Signal Integrity Checks: Detecting Manipulated and Linked Media (~2%)

Visual analysis detects subtle manipulation, artifacts, or unnatural capture paths that may look "clean" to the human eye. When combined with liveness, deduplication, and behavioural signals, image forensics help prevent visually convincing but fraudulent inputs from passing unchecked.

Ownership & Reuse Detection: Verification + Deduplication (~24%)

Face verification confirms the claimed identity matches the presented biometric. Deduplication links attempts across sessions, accounts, and platforms to expose organized identity reuse—catching syndicates that appear legitimate in isolation but reveal coordinated patterns over time.

Liveness: Active + Passive Detection (~73%)

Liveness detection stops the majority of suspected fraud because scalable attacks—replay videos, photo spoofs, screen presentations—fail here. Active and passive liveness together drive approximately 73% of flagged signals, making this the primary defence against automated biometric attacks.

This layered approach ensures that fraud caught at one layer doesn't overwhelm the next. High-volume, low-sophistication attacks are blocked early through liveness and entry controls. More sophisticated attempts—injection attacks, high-fidelity deepfakes, coordinated reuse—are surfaced through signal integrity checks and network intelligence. The result is a system that scales to handle millions of verification attempts while maintaining strong detection accuracy and low false positive rates.

From Verification to Network Intelligence

These verification layers form the foundation, but they are most effective when connected to broader intelligence across the customer lifecycle. Each verification decision generates signals—about the user, the device, the behaviour, the capture environment.

When these signals are isolated, they reveal only what happened in that single moment. When connected over time and across platforms, they reveal patterns: identity reuse, coordinated attacks, automation, insider assistance.

The diagram below illustrates how modern fraud defence operates across four interconnected zones, transforming individual verification events into continuous security infrastructure:

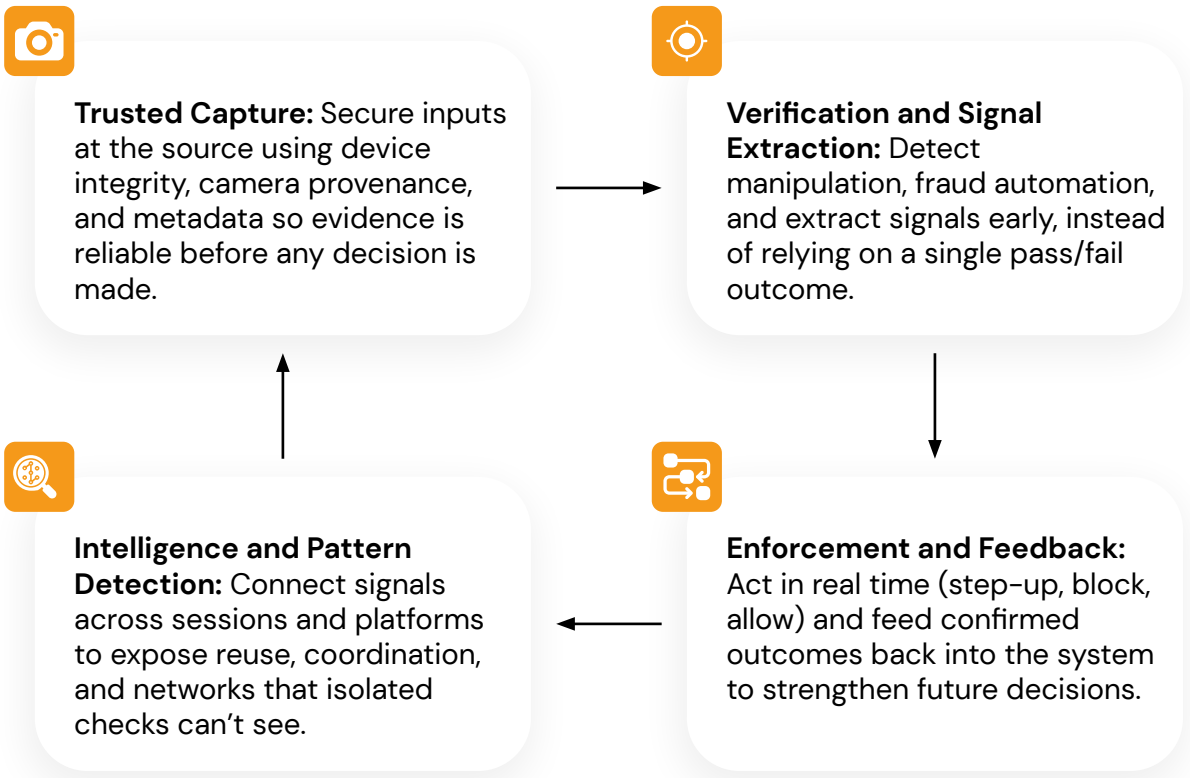

The three strategic priorities that follow—lifecycle intelligence, hardened authentication, and trusted capture—build on this foundation. They represent how institutions must apply these capabilities across the full identity journey to counter fraud that is now continuous, coordinated, and industrial in scale.

Smile Risk Intelligence

High-Level architecture overview.

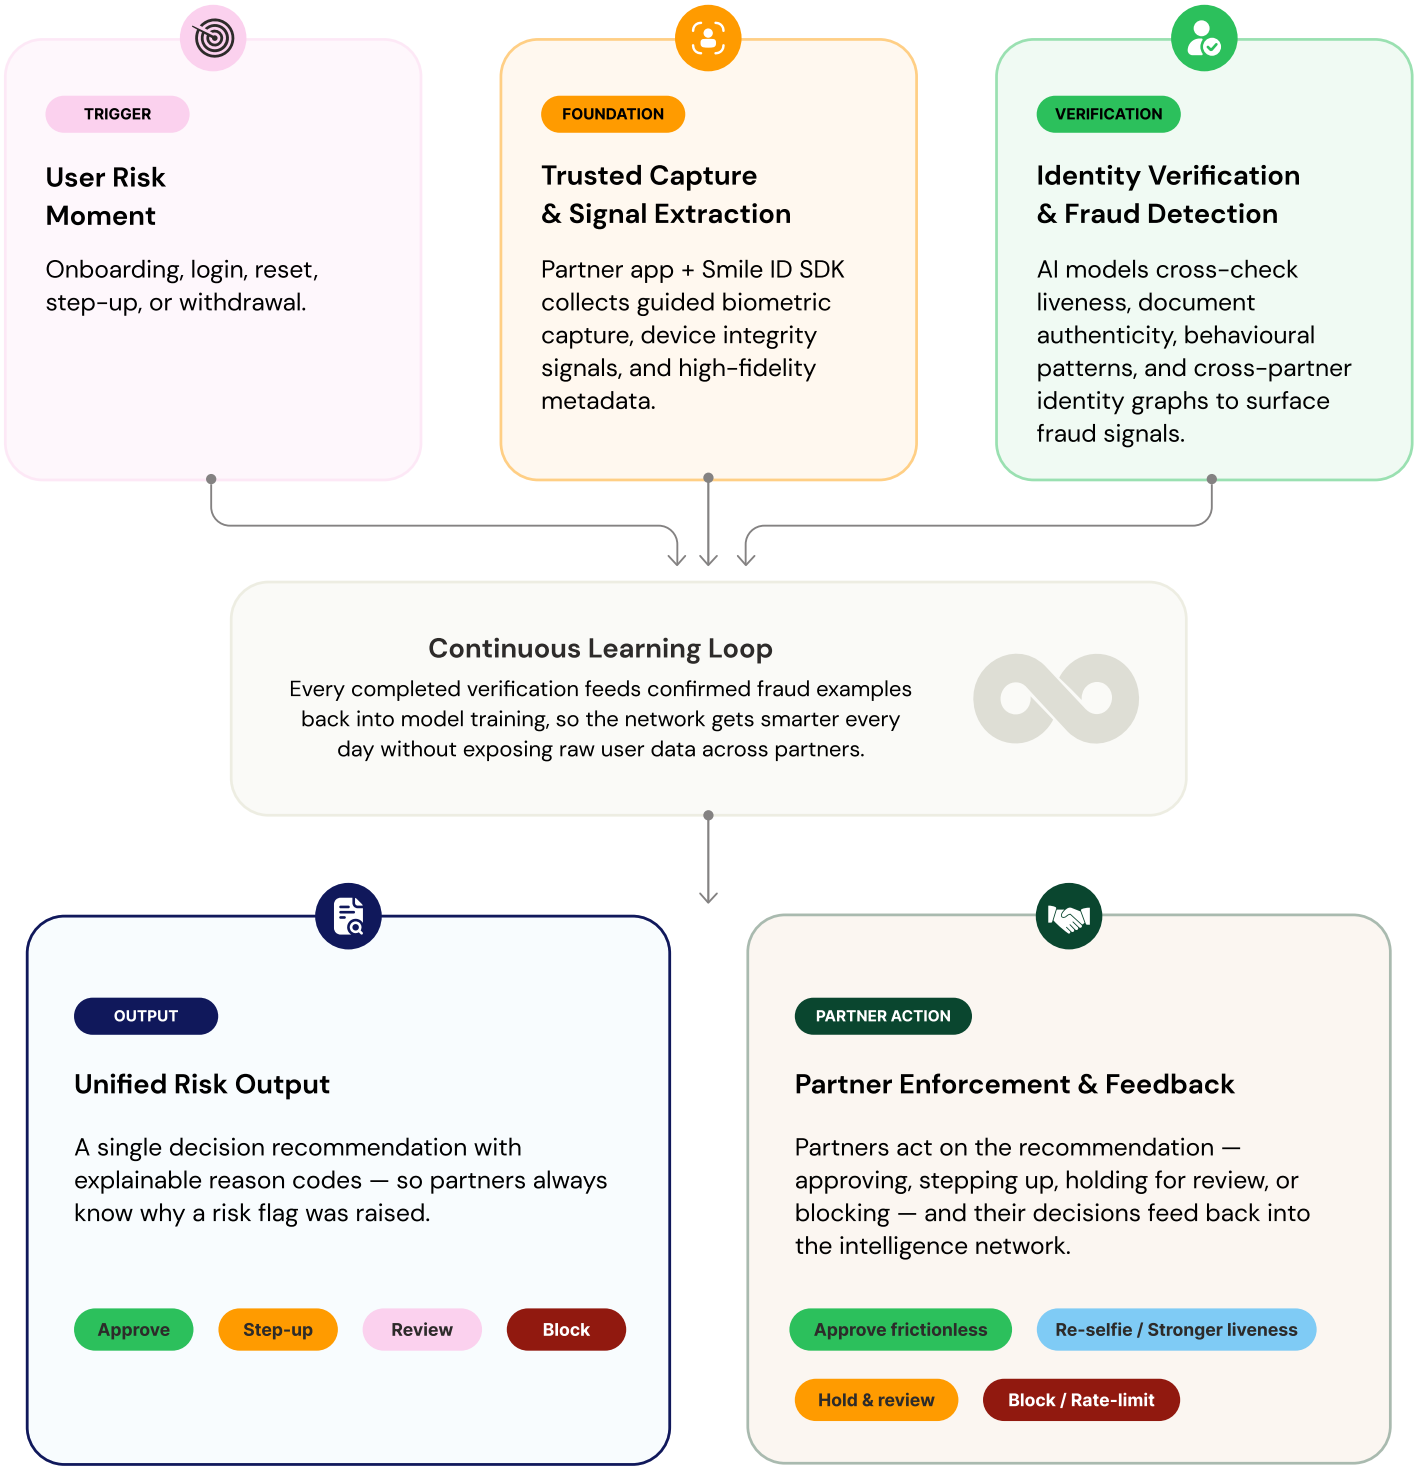

3.1 Priority One: Build Fraud Intelligence That Operates Across the Full Lifecycle

One-time verification is no longer sufficient. Attackers reuse the same identities, devices, and behavioural patterns across many sessions and services. A face that appears legitimate in isolation may be part of a syndicate operating at scale. A device that seems normal in one session may reveal automation patterns when viewed over time. A verification that passes today may be linked to confirmed fraud tomorrow.

Effective defence requires visibility across attempts and over time—connecting biometric, device, and behavioural signals so coordinated reuse and syndicate patterns surface early, before losses accumulate.

Why Lifecycle Intelligence Matters

In 2025, Smile Secure—Smile ID's biometric deduplication capability—detected 71% more duplicate fraud attempts than the combined total of 2023 and 2024. The increase reflects not just higher attack volume, but more sophisticated reuse: syndicates now operate identity supply chains, "aging" accounts through dormancy or low-risk activity before activating them for fraud or money laundering.

What Lifecycle Intelligence Enables

Fraud intelligence that operates across the lifecycle does three things isolated checkpoints cannot:

1. Detects identity reuse across accounts and platforms

The same face, document, or device appearing repeatedly becomes a liability rather than an advantage. Systems can flag when a "new" identity has already been linked to suspicious activity elsewhere in the network.

2. Surfaces behavioural patterns that indicate coordination

High-velocity attempts, burst patterns, shared devices or networks, and metadata consistency across supposedly unrelated users reveal syndicate operations that individual checks miss.

3. Improves enforcement over time through feedback loops

Each verification decision contributes intelligence back into the system. Confirmed fraud outcomes strengthen pattern detection. Approved users build trust signals that reduce friction. The system learns and adapts, rather than treating every interaction as independent.

From Verification Events to Network Intelligence

This approach requires moving from event-based thinking to network-based thinking. Instead of asking "is this attempt legitimate?" in isolation, systems must ask "how does this attempt connect to other attempts, across time and platforms?"

That shift has practical implications:

- Biometric matching becomes biometric tracking: Not just "does this face match the claimed identity?" but "has this face been seen before, and in what contexts?"
- Device signals become device history: Not just "is this device secure?" but "has this device been linked to suspicious patterns, shared identities, or high-velocity abuse?"
- Behavioural signals become behavioural graphs: Not just "is this action normal?" but "does this behaviour cluster with known fraud patterns across the network?"

This is networked fraud intelligence: systems that recognize repetition, reuse, and coordination—not just in single transactions, but across the full identity lifecycle.

Smile Secure: Network-Level Deduplication

Smile Secure detects when the same person attempts to register multiple times, even if their details or documents change. It matches applicants' biometrics against your user database and flags repeats, reducing multi-account abuse without adding friction for legitimate users. This capability is strengthened through Smile ID's partnership with Mastercard, supporting broader intelligence sharing and ecosystem-level protection where permitted under local data protection rules and Smile ID terms of service.

3.2 Priority Two: Harden Authentication at High-Value Moments

Fraud has moved decisively post-onboarding. Attackers no longer focus primarily on creating new accounts. They target verified users at the moments that unlock value: login and re-authentication, password resets and account recovery, device changes, and high-value transactions like withdrawals or limit increases.

These moments are structurally weaker than onboarding. Many platforms apply lighter controls at authentication because they assume the user has already been verified. Password resets rely on email links or SMS OTPs that can be intercepted or co-opted through social engineering. Device changes may require minimal proof of ownership. High-value withdrawals may trigger alerts but not always step-up verification. Attackers exploit these gaps, often with insider assistance, to gain control of accounts that have already passed KYC and built transaction history.

Why Post-Onboarding is the Battleground

The shift toward authentication fraud is visible across multiple data points. In West Africa, potential fraud attempts at retail banks rose approximately 50% year-over-year in 2025, driven primarily by verification volumes in authentication and account-recovery flows. The accounts being attacked were not new. They were verified, trusted, and often high value—making them more attractive targets than freshly created accounts that may face transaction limits or enhanced monitoring.

This pattern reflects how professional fraud syndicates operate. Instead of attempting to bypass onboarding repeatedly—which is noisy, friction-heavy, and increasingly well-defended—they focus on taking over accounts that have already been verified. Once inside, they can move funds quickly, exploit higher transaction limits, and cash out before controls react. The window for detection between takeover and exploitation is measured in minutes, not hours or days.

What High-Value Moments Are

Not all post-onboarding interactions carry the same risk. Effective defence requires identifying the moments where account control changes hands or value becomes accessible, and applying proportionate controls at those points:

1. **Login and re-authentication:** The first opportunity to access an account. If credentials are compromised—through phishing, SIM swap, or insider leaks—attackers can gain full account control.
2. **Password and credential resets:** Often the weakest link. Many platforms allow resets via email or SMS without requiring biometric verification, creating an easy path for attackers who have compromised communication channels.
3. **Device changes:** When a user switches devices or logs in from a new location, it may signal legitimate behaviour—or account takeover. Without step-up verification, attackers can migrate an account to their own infrastructure.
4. **High-value transactions:** Withdrawals, limit increases, large transfers, or admin changes (beneficiary updates, linked account modifications) concentrate value and should trigger stronger controls.

What Hardening Means in Practice

Hardening authentication does not mean applying maximum friction everywhere. It means applying friction proportionate to the risk and value at stake. For low-risk interactions; routine logins from known devices, small transactions within normal patterns—minimal friction maintains user experience.

For high-risk moments; first login from a new device, large withdrawal after password reset, adding a new payment beneficiary or admin changes following unusual activity—step-up verification is justified.

Practical hardening measures include:

1. **Multi-factor authentication at critical moments:** Require biometric verification in addition to OTP for password resets, device changes, and high-value transactions. Biometrics are harder to compromise than knowledge-based factors and provide proof of presence.
2. **Behavioural anomaly detection:** Flag logins or transactions that deviate from established patterns—unusual locations, high velocity, new transaction types or beneficiaries, high amounts—and require additional verification before proceeding.
3. **Time-based controls:** Introduce mandatory delays or cooling-off periods for high-risk actions initiated shortly after account access or credential changes, extending the opportunity for detection before attackers can cash out.
4. **Admin access restrictions:** Limit who can modify authentication settings, recovery methods, or linked accounts. Require senior approval or multi-party authorization for sensitive changes.

The goal is not to block users. It is to raise the cost and complexity of takeover attempts so that attackers cannot move quickly and quietly through post-onboarding flows.

Enhanced SmartSelfie™ for Step-Up Verification

Smile Secure detects when the same person attempts to register multiple times, even if their details or documents change. It matches applicants' biometrics against your user database and flags repeats, reducing multi-account abuse without adding friction for legitimate users. This capability is strengthened through Smile ID's partnership with Mastercard, supporting broader intelligence sharing and ecosystem-level protection where permitted under local data protection rules and Smile ID terms of service.

3.3 Priority Three: Protect Capture & Device Integrity at Source

Fraud detection is only as strong as the evidence it reviews. When attackers can compromise inputs—through virtual cameras, emulators, tampered apps, or injected media streams—the verification process is undermined before analysis even begins. A selfie can look perfect, pass liveness checks, and match a document, yet still be synthetic if the capture pipeline has been manipulated.

In 2025, injection-style attacks shifted from rare, sophisticated techniques to scalable fraud infrastructure deployed against African fintechs and banks. Smile ID's clustering analysis flagged more than 100,000 injection-style attempts per month, often linked to emulator farms, virtual cameras, or SDK tampering. These attacks succeed not by "beating the models" but by bypassing capture entirely—feeding pre-recorded or AI-generated media into verification systems as if it came from a live camera on a real device.

Why Capture Integrity is Foundational

Traditional fraud detection operates on the assumption that the media being reviewed—a selfie, a document photo—was captured in a legitimate environment. If that assumption is false, downstream checks become unreliable. Visual analysis may detect obvious manipulation, but high-quality injection attacks produce media that looks indistinguishable from genuine capture.

The fraud is revealed not through the image itself, but through signals indicating how the image was produced: metadata inconsistencies, abnormal capture timing, conflicting hardware fingerprints, or environmental cues that suggest virtualization or automation.

This is why nearly 90% of verifications rejected for suspected fraud in 2025 were caught using SDK-based integrations—up from 15% in 2023 and 65% in 2024. SDKs sit inside the application and can help validate that media came from a live camera on a real device in a trusted environment, not from a virtual camera, emulator, or tampered client.

API-based implementations see the final output but have limited visibility into how it was created, leaving gaps that injection attacks exploit.

Fraud Caught by Integration

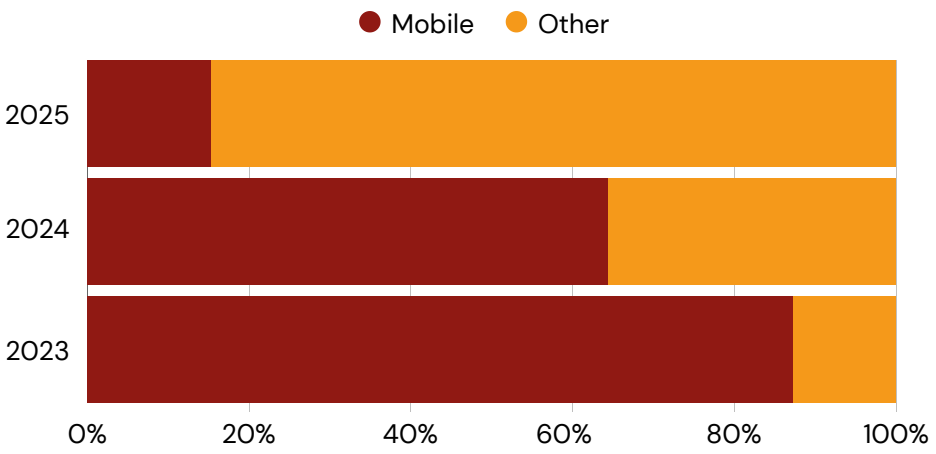

API Integration — Unknown Capture Source

Capture conditions are unverified; metadata may be absent, inconsistent, or fabricated.

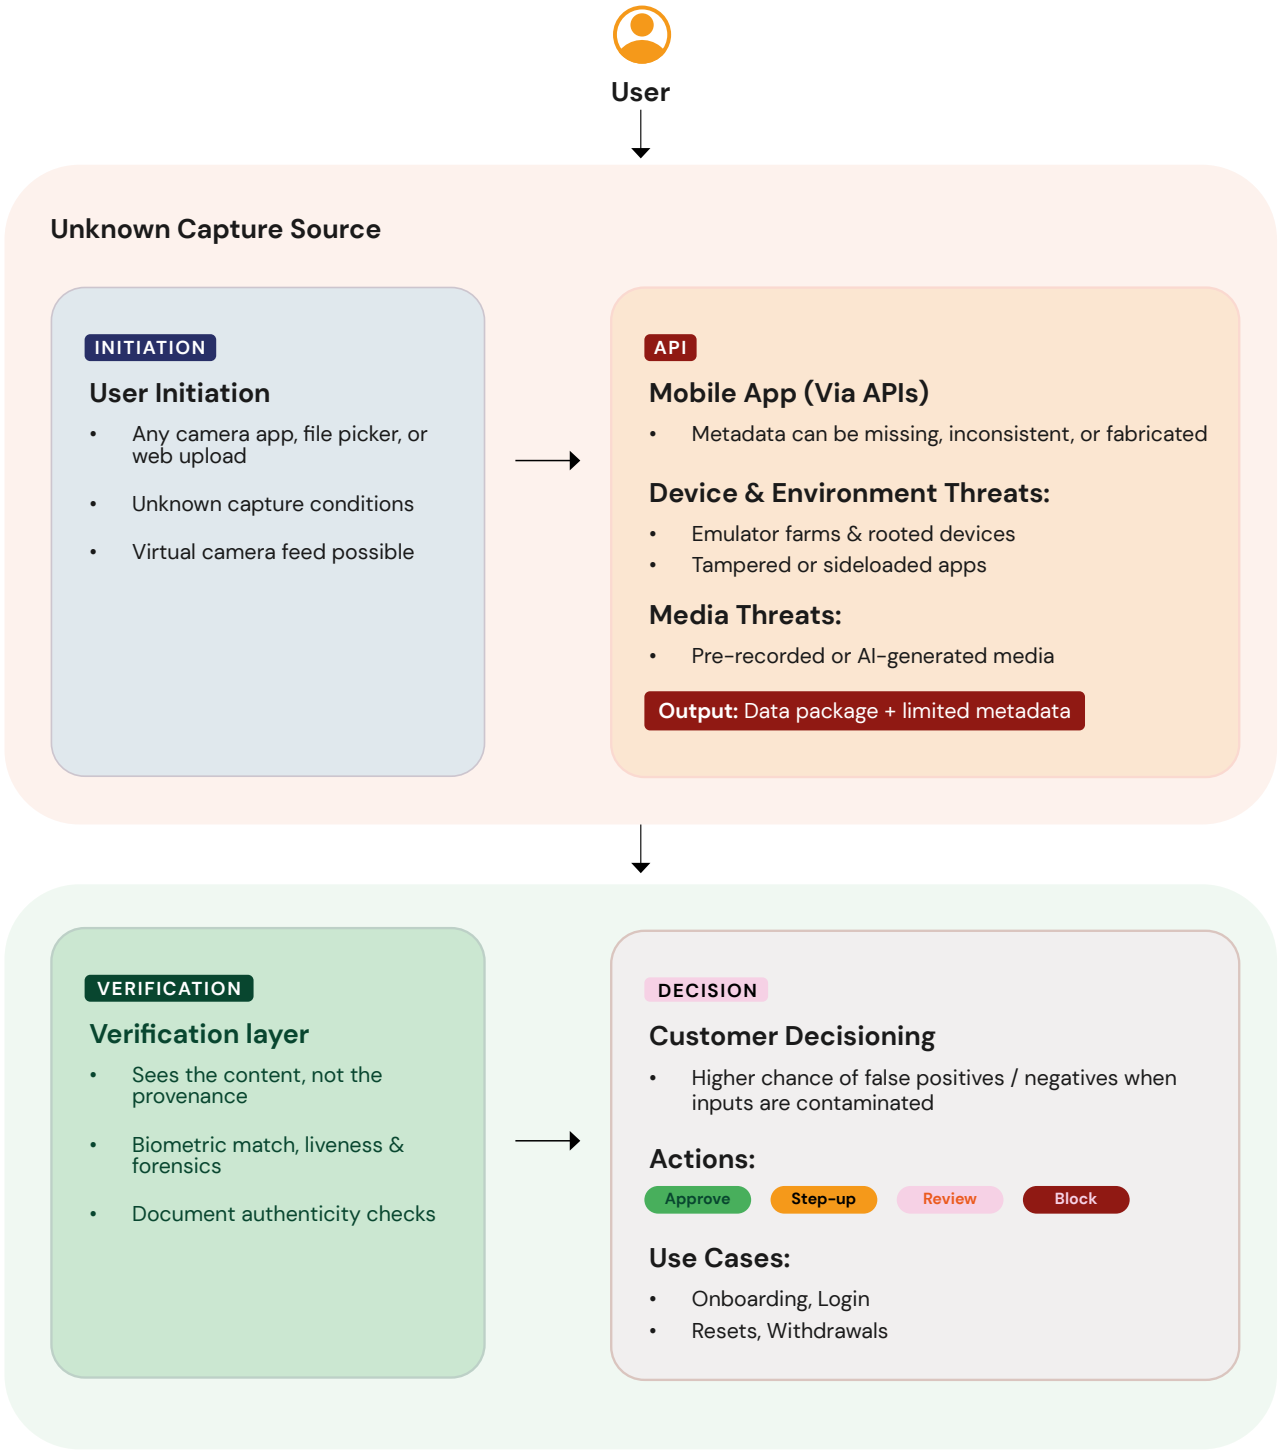

High Risk of Injection Attacks via API integration

In API flows, systems see the data but have limited visibility into how it was produced. In capture-controlled flows, integrity is enforced at creation, so downstream checks can rely on stronger, tamper-resistant signals.

SDK Integration — Integrity at Capture

Harder to tamper with; richer fraud signals available at the point of capture.

User

Real-World User Interaction

CAPTURE

SDK Capture Module

- Guided capture & quality checks
- Active liveness prompts (where configured)

INTEGRITY GATE

On-Device Integrity Gate

Device Posture:
Rooted / jailbroken, emulator / virtual env, debugger, overlay

App Integrity:
Signature mismatch, sideloaded APK, tampered app

Camera Provenance:
Virtual camera detection / sensor consistency

Automation Signals:
Burst attempts, scripted patterns

PAYLOAD

Payload Signing

- Selfie + doc + liveness data
- Device fingerprint + OS / SDK version + capture telemetry
- Network signals: IP, VPN / proxy hints

Verification & Fraud Intelligence Layer

BIOMETRIC

Biometric Verification:
Face match + liveness + forensics

DOCUMENT

Document Verification:
Authenticity, tamper & data checks

COMPLIANCE

AML Check
Watchlist & sanctions screening

INTELLIGENCE

Fraud Intelligence

Graph / Cluster Detection:
- Reused faces, devices, or PII;
- Abnormal velocity or patterns

Shared Signals:
- Confirmed fraud feedback;
- Blocklists & clusters

Device Metadata:
- Cross-referenced from payload

DECISION

Customer Decisioning

Actions:

Approve | Step-up | Review | Block

Use Cases:
- Onboarding, Login
- Resets, Withdrawals

What Trusted Capture Requires

Protecting capture integrity means validating three things: the device environment, the capture process, and the delivery pathway.

1. **Device environment validation:** Is the device real or emulated? Has the operating system been modified (rooted/jailbroken)? Are security protections intact, or has the app been tampered with? Device-level signals—hardware fingerprints, OS integrity checks, and environment assessments—help distinguish genuine devices from virtualized fraud infrastructure.

2. **Capture process integrity:** Did the media come from the device's physical camera, or was it injected through a virtual camera or pre-loaded stream? Are capture timing and metadata consistent with human behaviour, or do they suggest automation? Real-time validation during capture detects substitution attempts before media reaches the server.

3. **Delivery pathway security:** Was the media transmitted securely, or could it have been intercepted and modified in transit? Are the credentials and endpoints legitimate, or are they routing through proxies or unauthorized services? Validating the full pipeline—from capture to submission—reduces the risk of payload manipulation.

How Capture Integrity Enables Downstream Detection

Trusted capture does not stop fraud by itself. But it makes every other layer stronger. When systems can rely on capture integrity, they can:

Reduce false positives: Clean, validated inputs mean fewer ambiguous cases where genuine users are flagged due to poor image quality or environmental noise.

Surface injection patterns earlier: Clustering analysis can link attempts based on capture anomalies—shared device fingerprints, repeated metadata signatures, or behavioural consistencies—that only become visible when capture signals are available.

Support network intelligence: Device and behavioural signals feed into broader pattern detection, helping identify syndicates that reuse infrastructure across platforms.

Improve enforcement confidence: When capture is trusted, systems can act more decisively on suspicious signals without fearing they are rejecting legitimate users due to noisy data.

In short: capture integrity turns identity verification from a visual inspection process into a systems validation process. It shifts the question from "does this look real?" to "can we trust where this came from?"

SDK Controls for Secure Capture

Smile ID's mobile SDKs validate capture integrity at the source by confirming that media originates from a live camera on a real device in a trusted environment—not from virtual cameras, emulators, or tampered clients. Because the SDK operates inside the application, it can check device posture, enforce secure capture, and collect richer device, network, and behavioural signals that API-only flows cannot access. This layered approach makes injection attacks significantly harder and strengthens Smile ID's ability to detect coordinated fraud through networked intelligence.

3.4 Practical Implementation: Designing for Emerging Market Conditions

The three priorities outlined above—lifecycle intelligence, hardened authentication, and trusted capture—form the foundation of effective defence. But implementing them across emerging digital economies requires acknowledging real-world constraints: intermittent connectivity, variable device quality, assisted onboarding workflows, and diverse user populations with different levels of digital literacy.

These constraints are not edge cases. They define how millions of users access financial services. Platforms that ignore them risk excluding legitimate users. But platforms that treat all use cases as equally constrained risk applying insufficient controls to high-value accounts and transactions, creating exploitable gaps for professional fraudsters.

The Trade-Off: Inclusion vs. Security

There is a legitimate tension between reaching more people and applying stronger controls. Assisted capture, offline verification, WhatsApp and progressive web applications, USSD-based authentication, and agent-mediated onboarding expand access, but they also widen the gap between evidence capture and validation, increase reliance on documents over biometrics, on APIs over SDKs, and reduce visibility into device and behavioural signals.

This trade-off is acceptable for certain use cases. Onboarding new users in low-connectivity environments, enabling small-value transactions for underbanked populations, and supporting agent-assisted flows in rural markets—all justify lighter controls in exchange for broader inclusion.

The risk is manageable when transaction limits are low, funds are not immediately liquid, and account activity can be monitored over time.

But the same implementation cannot serve all use cases. When value at risk rises—high-balance accounts, large withdrawals, access to liquid international markets like crypto or cross-border remittances—controls must tighten proportionately.

Platforms that apply the same lightweight verification to mass-market onboarding and high-value authentication are structurally vulnerable to targeted takeover attacks.

Where to Tighten Controls

Effective implementation requires segmenting use cases by risk and applying controls accordingly:

Mass-market onboarding (low transaction limits, monitored activity):

Assisted capture, document-led verification, and hybrid workflows are acceptable. The goal is inclusion, with fraud risk managed through transaction limits, velocity controls, and behavioural monitoring over time.

High-value account access (authentication, recovery, admin changes):

Require SDK-based, biometric verification with dynamic liveness, enforce device integrity checks, and apply step-up controls with multi-factor authentication at critical moments. The goal is preventing takeover of accounts that can move significant value quickly.

Liquid, high-risk markets (crypto, remittances, investment platforms):

Enforce the strongest controls at onboarding and throughout the lifecycle. These platforms attract professional fraud syndicates because funds can be moved across borders and cashed out rapidly. Even "low-value" accounts can serve as mule infrastructure for laundering.

The key insight:

Implementation can vary by value at risk.. Platforms can support both inclusion and security by tiering controls—applying lighter verification where risk is low and value is contained, and hardening defences where risk and liquidity are high.

Channel Mix and Infrastructure Constraints

Some channels inherently limit what controls can be applied. USSD-based services cannot enforce biometric verification. WhatsApp-based onboarding cannot validate device integrity. Web-based flows lack the same device-level signals that mobile SDKs provide.

These channels serve important roles in expanding access, but they should not be the only—or even primary—path for high-value interactions.

Platforms operating under infrastructure constraints should:

Prioritise SDK-based mobile apps for high-value flows: Where possible, migrate authentication, recovery, and high-risk transactions to app-based channels that support trusted capture and device validation.

Apply compensating controls in data constrained channels: If USSD or web-based verification is necessary, layer in additional security—transaction limits, eKYC checks, velocity monitoring, mandatory cooling-off periods, or agent verification—to offset weaker identity signals.

Limit functionality in assisted or deferred workflows: Where capture happens offline or through agents, restrict immediate access to high-value features. Require step-up verification before enabling withdrawals, limit increases, or sensitive account changes.

Minimize the time between capture and validation: Longer delays expand the substitution window. Where possible, enforce synchronous verification or near real-time submission to reduce opportunities for media swapping or replay.

Building for Scale Without Compromising Security

As platforms grow, the operational challenge is maintaining strong controls without creating friction that drives users away. The solution is not to weaken controls uniformly, but to apply them intelligently:

- **Use risk-based decisioning:** Apply friction only where signals indicate elevated risk. Routine interactions from known devices and locations can remain frictionless.
- **Invest in user education:** Many users will accept step-up verification if they understand it protects their funds. Clear communication reduces abandonment. Tell users which things you will never ask them to do, to help inoculate them against social engineering.
- **Monitor control effectiveness over time:** Track false positive rates, user drop-off, and fraud outcomes to identify where controls are too tight or too loose, and adjust accordingly.

The goal is not to choose between inclusion and security. It is to apply the right level of security to the right use cases, so that platforms can serve broad populations while protecting high-value accounts from professional fraud syndicates.

Conclusion

From Selfies to Signals: Trust in the Security Era

The 2025 patterns in this report point to a clear shift: the biggest identity losses are not driven by "better fakes" alone. They are driven by control over the pipeline; how identity evidence is captured, what environment it comes from, how it is reused, and how it is exploited after approval.

Three forces are converging to reshape fraud and defence:

Fraud is cheaper and faster to run.

AI and automation have lowered the cost of producing convincing media and running iterative, high-volume attempts—making repeat attacks the norm.

Fraud is built around data reuse.

Syndicates don't need a new technique for every target. They reuse faces, devices, and playbooks across platforms—finding the seams where controls are lighter: login, account recovery, SIM swaps, device changes, and withdrawals.

Fraud increasingly wins through contamination.

Attackers tamper with inputs—emulators, virtual cameras, modified apps, proxies, forged metadata—so systems assess evidence that may not come from a live capture.

Regulation is moving the same way. In 2025, AML/CFT expectations continued expanding beyond banks to fintechs, wallets, agents, remittance providers, and merchants connected to regulated networks. Compliance is no longer a bank-only concern—it is an ecosystem-wide requirement.

What Effective Defence Requires

The three strategic priorities outlined in this report—lifecycle intelligence, hardened authentication, and trusted capture—are interconnected, not independent. Capture integrity enables richer signals.

Lifecycle intelligence reveals where fraud concentrates. Hardened authentication applies proportionate friction at those moments. Together, they shift defence from isolated checkpoints to continuous security infrastructure.

Platforms that treat identity as a one-time verification event will continue losing accounts to takeover. Platforms that connect signals across the lifecycle, harden controls where value is accessible, and validate evidence at its source will turn attackers' repetition into a liability.

Network Defence: Trust in the Security Era

As automation lowers the cost of fraud and makes it easier to repeat at scale, advantage won't come from any single model, check, or moment.

It comes from securing inputs, linking signals over time, and applying proportionate controls where value moves—so repetition becomes a liability, not an edge.

Fraud now operates as repeatable, networked infrastructure. Defence must do the same. This approach—a Network Defence—connects signals across the identity lifecycle, detects coordination that isolated systems miss, and strengthens with every verification.

In the security era, trust isn't granted once. It is earned, defended, and re-earned—continuously. Identity is no longer a checkpoint. It is security infrastructure.

Glossary (the A-Z of Fraud)

Fraud Types & Attack Methods

Account Takeover (ATO): The unauthorised access to a legitimate user's account, typically after onboarding, often used to move funds or exploit trusted access.

Authentication Fraud: Fraud that occurs during login, re-authentication, or account access rather than at initial sign-up.

Bonus / Referral Abuse: The exploitation of promotional incentives by creating or reusing identities to claim rewards multiple times.

Chargeback Fraud: The act of disputing legitimate transactions after goods or services have been received causes financial loss to merchants.

Deepfake: AI-generated or manipulated images or videos designed to impersonate a real person during biometric verification.

Document Fraud: The use of forged, altered, stolen, or AI-generated identity documents to misrepresent a user's identity.

Face Swap: A biometric attack where one person's facial features are digitally overlaid onto another's image or video to bypass identity checks.

Friendly Fraud: A form of chargeback fraud where a user falsely claims a legitimate transaction was unauthorised.

Identity Farming: The large-scale collection or creation of identities—real or synthetic—for repeated fraudulent use across platforms.

Injection Attack: A technique where pre-recorded or synthetic media is fed directly into a verification system, bypassing the physical camera.

Insider-Assisted Fraud: Fraud enabled or facilitated by individuals within an organisation who abuse legitimate access or systems.

Money Laundering (Layering): The movement of illicit funds through multiple accounts or platforms to obscure their origin and evade detection.

Presentation Attack: An attempt to fool biometric systems using physical or digital representations, such as photos, videos, or masks.

Replay Attack: The reuse of previously captured biometric or session data to impersonate a legitimate user.

Synthetic Identity: An identity constructed using a combination of real and fabricated information that does not correspond to a single real person.

Virtual Camera Attack: An attack that replaces a device's real camera feed with a fake or preloaded stream during identity verification.

Identity, Biometrics & Verification Concepts

Active Liveness: A biometric technique that requires user interaction, such as movement or gestures, to confirm a real person is present during verification.

Authentication: The process of confirming a user's identity when they attempt to access an existing account or perform a sensitive action.

Biometric Verification: The use of physical characteristics, such as a face, to confirm that a person is who they claim to be.

Deduplication: The detection of repeated or reused identities by identifying the same biometric or identity data appearing across multiple accounts.

Dynamic Liveness Detection: An advanced liveness method that adapts in real time to user behaviour and environmental conditions to resist spoofing and deepfake attacks.

Face Embedding: A mathematical representation of a face used by biometric systems to compare and identify facial similarity across checks.

Facial Biometric: Identity verification based on facial features, commonly used in selfie-based onboarding and authentication flows.

Identity Lifecycle: The full journey of an identity from onboarding through authentication, account changes, and ongoing use over time.

Liveness Detection: Techniques used to determine whether a biometric sample comes from a live person rather than a photo, video, or synthetic source.

Passive Liveness: A liveness approach that evaluates biometric authenticity without requiring explicit user actions.

Re-authentication: The process of verifying a user's identity again after onboarding, often during login or high-risk actions.

SmartSelfie™: Smile ID's facial biometric verification technology used to confirm identity during onboarding and authentication.

Device, Network & Infrastructure Signals

Device Fingerprint: A composite identifier created from device attributes that helps detect repeated or abnormal usage across identity checks.

Device Integrity: An assessment of whether a device has been altered, compromised, or is operating in an unsafe environment.

Emulator: A simulated device environment that allows fraudsters to run multiple identity checks without using physical phones.

Injection Vector: The technical pathway used to feed fake or manipulated data into an identity verification system.

Jailbroken Device: An iOS device that has been modified to bypass manufacturer security restrictions, increasing fraud risk.

Metadata: Hidden technical information attached to images, videos, or sessions that provides context about how data was captured.

Network Signals: Indicators related to connectivity, such as IP address behaviour, VPN usage, or proxy routing, used to assess risk.

Rooted Device: An Android device that has been modified to grant elevated access, often enabling fraud tools or tampering.

SDK (Software Development Kit): A packaged set of tools embedded in an app that enables secure identity capture and fraud signal collection.

Tampered Application: An app that has been altered or modified to bypass security controls or manipulate identity data.

Virtual Environment: A non-physical operating environment, such as a virtual machine or container, commonly used in automated fraud.

VPN / Proxy Detection: The identification of traffic routed through anonymising services that can obscure a user's true location or identity.

Behavioural, Pattern & Network Intelligence

Behavioural Analysis: The evaluation of user actions over time to identify abnormal or suspicious activity that may indicate fraud.

Behavioural Graph: A network model that maps relationships between identities, devices, and actions to surface coordinated fraud patterns.

Cross-Platform Fraud: Fraud activity that spans multiple products, platforms, or institutions rather than targeting a single system.

Face Clustering: The grouping of similar facial biometrics to detect repeated or reused identities across multiple accounts.

Fraud Syndicates: A coordinated group of actors using shared tools, identities, or infrastructure to commit fraud at scale.

Identity Graph: A connected representation of identities and their shared attributes, revealing relationships that individual checks cannot detect.

Identity Reuse: The repeated use of the same identity elements, such as faces or documents, across different accounts or platforms.

Network Intelligence: Fraud detection insights derived from analysing patterns and signals across a broader ecosystem rather than a single platform.

Network Defence: A security approach that connects fraud signals across the identity lifecycle and across multiple platforms, not just within a single institution. By aggregating intelligence from verification attempts across clients, Network Defence detects coordination, identity reuse, and syndicate patterns that isolated systems miss—strengthening defences over time by treating fraud detection as continuous, shared infrastructure rather than isolated checkpoints.

Pattern Detection: The identification of repeated behaviours or signals that indicate organised or automated fraud.

Risk Scoring: The assignment of a dynamic risk level to an identity or transaction based on multiple signals and behaviours.

Trust Mesh: A shared framework where identity and fraud signals are collectively analysed to strengthen ecosystem-wide security.

Regulatory, Policy & Compliance Terms

AML (Anti-Money Laundering): A set of laws, controls, and processes designed to prevent illicit funds from entering the financial system.

CFT (Counter Financing of Terrorism): Regulatory measures aimed at detecting and preventing the flow of funds to terrorist organisations.

Digital ID: A government-issued or government-recognised identity that can be verified electronically.

FATF (Financial Action Task Force): An international body that sets global standards for combating money laundering and terrorist financing.

KYC (Know Your Customer): The process of verifying a customer's identity to meet regulatory and risk management requirements.

PEP (Politically Exposed Person): An individual with a prominent public role who is exposed to a higher financial crime risk due to their position.

VASP (Virtual Asset Service Provider): An entity that facilitates the exchange, custody, or transfer of digital assets and is subject to financial regulation.

Smile ID Platform & Intelligence

Enhanced SmartSelfie™: An advanced facial biometric verification solution that combines liveness detection and fraud analysis to prevent spoofing and impersonation.

Fraud Intelligence: The analysis of identity, device, behavioural, and network signals to detect and prevent fraud across the user lifecycle.

Smile Secure: Smile ID's deduplication and identity reuse detection capability designed to surface duplicate or synthetic identities within and across platforms.

Smile ID Risk Intelligence: Smile ID's next-generation fraud framework that integrates device integrity, biometric analysis, behavioural patterns, known bad actors and network-level signals into a unified risk view.

Contributors

Nicholas Kanyagia
Data Analyst

Emmanuel Agwu
Marketing Manager

Mark Straub
CEO

Lameck Orina
Multimedia Designer

Gift Arku
Marketing Associate

Marie Kruger
Head of Marketing

Mala Goel
GTM Lead

Kasim Sodangi
Director, Governmental and Regulatory Affairs