MICROBLINK | EBOOK

Mapping the Rise of AI-Powered Identity Fraud

A data-driven look inside the global evolution of attacks

Methodology

This report is powered exclusively by proprietary data gathered during the first half of 2026 from our global scanning footprint, as well as a few outside data sources where noted. While others rely on sentiment surveys, we rely on **hard telemetry.** By analyzing millions of identity interactions across our ecosystem, we have uncovered a "Sophistication Paradox": as digital trust grows, the methods to subvert it become more surgical.

Note: *The following insights reflect trends observed within our internal traffic and confirmed fraud captures as of Q1 2026.*

Fraud Without Boundaries

Every day, there are millions of digital interactions across regions, industries, and use cases. When examining this data, it is clear that fraud is no longer uniform. It is becoming increasingly targeted, regionalized, and sophisticated.

Attack patterns vary significantly by market, shaped by local regulations, document types, and user behaviors. Fraudsters are adapting their tactics accordingly; testing, refining, and redeploying methods that work in one region into others with similar vulnerabilities. At the same time, advances in generative AI are accelerating this cycle, enabling faster iteration and higher-quality attacks with minimal effort.

The result is a more dynamic threat landscape, where fraud is not only harder to detect, but constantly evolving. Organizations can no longer rely on static defenses or one-time verification. Understanding and acting on how fraud behaves across different regions and channels is now critical to staying ahead.

Hartley Thompson III
CEO Microblink

The Many Aspects of Deepfakes

Deepfakes are no longer limited to manipulated videos. Advances in generative AI have made it possible to create highly realistic synthetic content across multiple modalities, including text, voice, images, documents, and video. More importantly, these attacks can now be generated and deployed at unprecedented scale. Just as phishing became profitable through automation and volume, AI is enabling fraudsters to industrialize identity fraud, social engineering, account creation, and impersonation attacks.

What was once limited by human effort can now be executed continuously across thousands of targets simultaneously. This shift is giving rise to a new era of agentic fraud, where AI is used not only to create realistic content, but also to automate fraud workflows, orchestrate attacks, and adapt interactions in real time. Understanding the different types of deepfakes is the first step toward building defenses against fraud that is becoming increasingly automated, scalable, and convincing.

Deepfake video

AI-generated or manipulated videos that replace or alter a person's face and movements to create realistic but false visual content.

AI-generated images

Synthetic images created from scratch that depict people, places, or events that never existed.

Synthetic identity

Fake identities constructed using a combination of real and fabricated information, often supported by AI-generated documents and images.

Deepfake avatars

AI-generated digital humans used in real-time interactions, customer support scams, or social engineering attacks.

Voice cloning

AI-generated speech that mimics a person's voice, including tone, emotion, and cadence, often used in phone scams and impersonation attacks.

Real-time deepfakes

Live manipulation of video or audio during calls and streams, enabling real-time impersonation.

Deepfakes Are the New Norm

Deepfakes and generative AI fraud are no longer emerging threats confined to sophisticated cybercriminal groups. They are rapidly becoming the new baseline for digital fraud.

This shift has fundamentally changed the economics of fraud. Generative AI has made fraud faster, cheaper, and easier to scale than ever before. They are quickly becoming the standard operating model for modern digital attacks.

$40B

in U.S. **fraud losses** due to Gen AI by 2027

30%

of enterprises will consider ID verification solutions **unreliable due to deepfakes**

700%

surge in deepfake document forgery and related **AI-driven identity fraud recorded** in early 2026 compared to 2024

Fraud Targets Both High-Value Economies and Obscure Regions

When we look at our internal document traffic, the "Fake Rate" isn't evenly distributed. As seen in the map here the Canadian documents see the highest fake rate, at 43% followed by the U.S. at 30%

However, the countries with the next highest fake rates are Madagascar at 22%, Puerto Rico at 19% and Malawi at 18%. This means, while fraud mostly targets high-level economies, it distributes fairly evenly across regions.

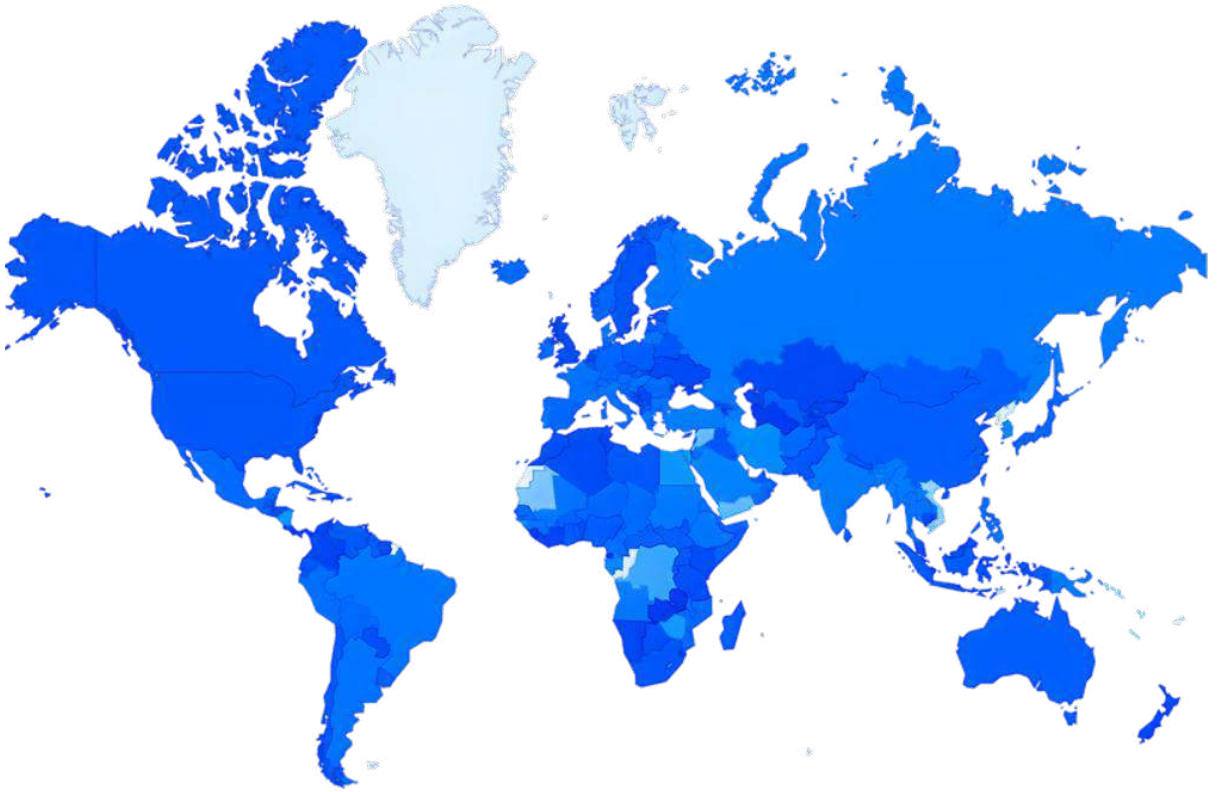

Why Fraud is Country-Agnostic

Fraud concentrates where the money is. That is why the U.S. and Canada show up as high-volume targets: large digital economies, mature credit systems, high-value financial accounts, and widely accepted identity credentials make them attractive to organized fraud operations.

The figures shown here are derived from Microblink's internal fraud datasets, which are intentionally enriched with fraudulent and manipulated identity documents to support fraud analysis and detection research. As a result, they reflect patterns observed within that analyzed traffic rather than the overall volume of identity verification transactions occurring within a country.

Within those datasets, Canada and the United States consistently appear among the most targeted markets. This reflects the value of those identities to fraudsters, who seek credentials that can be reused across financial services, marketplaces, government programs, and other digital ecosystems.

At the same time, attackers increasingly probe smaller and less obvious markets. Documents from these regions may be less familiar to fraud teams, less represented in historical fraud datasets, and less extensively tested across verification systems. In other words, the U.S. and Canada are targeted because they are lucrative. Smaller countries are targeted because they can present different verification challenges.

Country	Observed Fake Rate	Market Insight
CANADA	**43.48%**	Nearly 1 in 2 documents flagged in our Canadian traffic shows signs of manipulation.
USA	**30.73%**	A massive volume of sophisticated attacks targeting the world's largest economy.
MADAGASCAR	**22.00%**	Proves that high-sophistication fraud has no geographic boundaries.
PUERTO RICO	**19.33%**	High-pressure zone requiring specialized regional document knowledge.
MALAWI	**18.03%**	Emerging markets are seeing a surge in organized fraud.

THE ANATOMY OF AN ATTACK

Three Global Personas

Our Fraud Lab has categorized the failures we see into distinct "Attack Personas." Depending on where your users are, the threat changes completely. What looks like a single global fraud problem is really a shifting cast of characters, each adapting to local conditions like a chameleon changing colors mid-stride. In some regions, attackers lean heavily into fully digital deception, exploiting the gap between what a camera sees and what a human assumes is “real.” In others, the craft moves closer to the document itself, with precision edits that turn legitimate IDs into near-perfect impostors. And in many parts of the world, the playbook isn’t futuristic at all. It is grounded in physical reproduction techniques that exploit how often businesses still rely on visual inspection or low-fidelity capture.

The result is a fragmented threat landscape where fraud does not evolve uniformly. It specializes.

PROFILE #1

The ‘Digital Ghost’

PROFILE #2

The ‘Master Forger’

PROFILE #3

The ‘Analog Architect’

PROFILE #1

The ‘Digital Ghost’

In Europe and Asia, attackers have refined screen-based fraud into something closer to a performance than a trick. They present ultra-high-resolution images or synthetic IDs on phones, tablets, or monitors, carefully controlling brightness, refresh rates, and glare so the capture system “believes” it’s seeing a physical document. To the naked eye, it can look almost too clean, but without analysis of pixel-level artifacts like moiré patterns, backlight bleed, or unnatural edge sharpness, these attacks slip through as if they were the real thing.

What makes these attacks particularly effective is that the tools required are no longer limited to sophisticated fraud rings. High-resolution displays, AI-generated identity documents, image enhancement tools, and synthetic media applications are widely accessible, allowing attackers to create convincing presentations with minimal technical expertise.

The takeaway
If your provider isn't analyzing pixel-moiré and screen-glare, they are effectively blind in Europe.

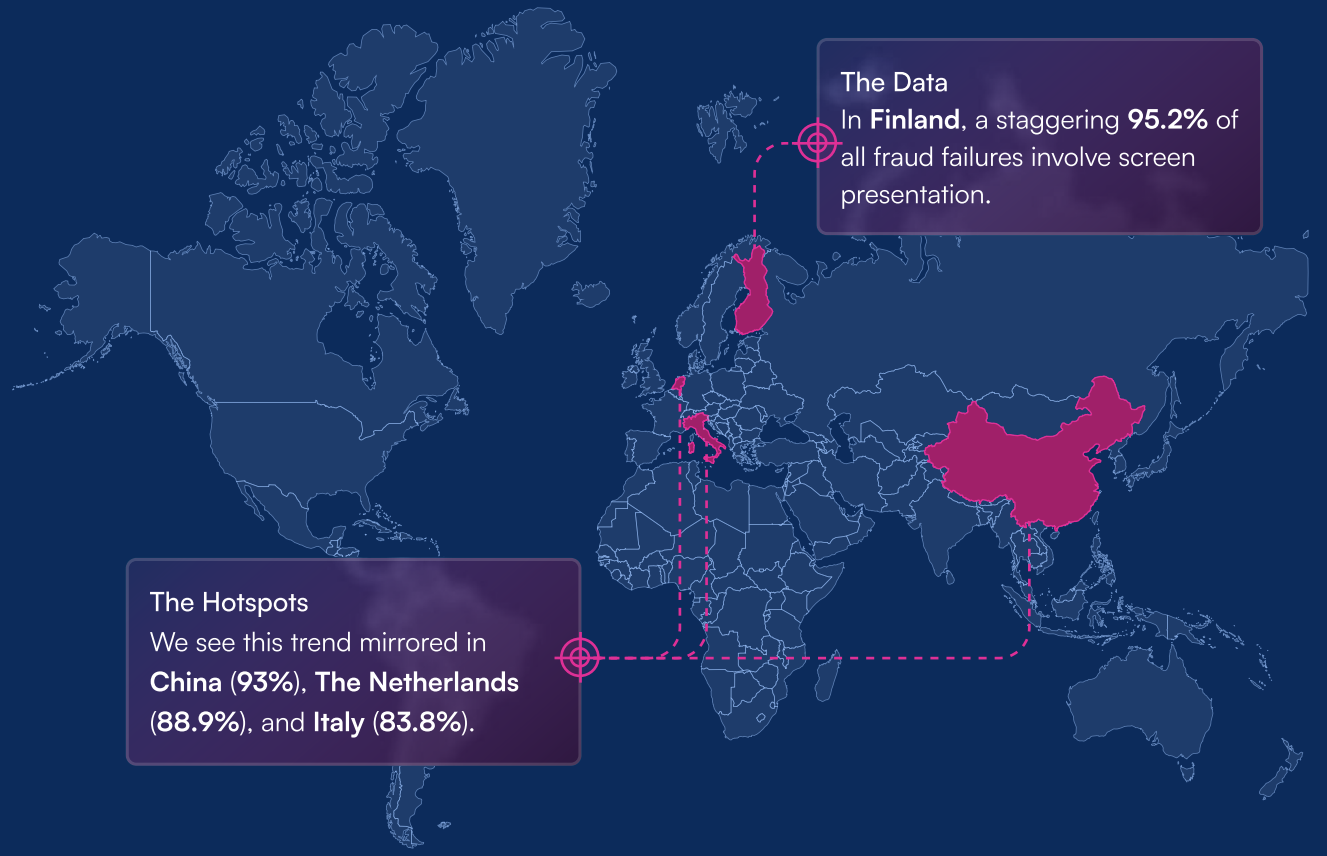

PROFILE #2

The ‘Master Forger’

In North America, the attack moves from presentation to precision. Fraudsters work directly on the document, swapping portraits, editing fields, or reconstructing elements so the ID still passes basic structural checks while quietly becoming someone else’s identity. These aren’t crude edits. They preserve barcodes, layouts, and security features just enough to appear legitimate, which means detection depends on catching subtle inconsistencies in image integrity, layering, and data relationships rather than obvious visual flaws.

AI-powered image editing tools, synthetic document generators, and widely available design software have dramatically lowered the expertise required to alter identity documents convincingly. Fraudsters can now modify legitimate credentials with a level of speed and realism that was difficult to achieve just a few years ago.

The takeaway
This is the "Sophistication Paradox." The more secure the document, the more sophisticated the forgery tools become.

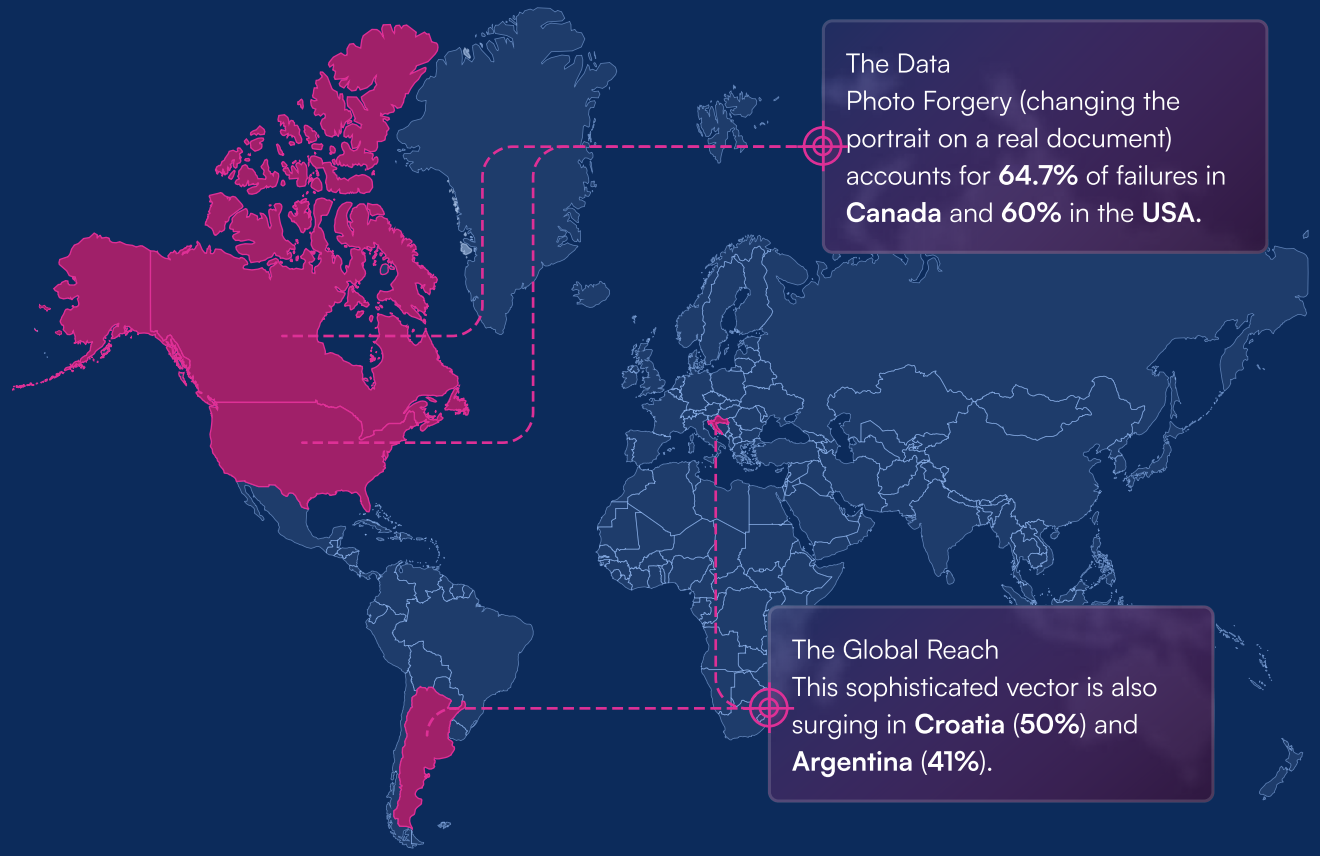

PROFILE #3

The Analog Architect

In many regions, “old school” still wins because it exploits a modern blind spot. High-resolution printers, inexpensive imaging equipment, and widely available editing tools have made it easier than ever to create convincing physical replicas of identity documents. Attackers no longer need sophisticated forgery operations to produce credentials that appear authentic during standard verification flows.

Attackers produce high-quality photocopies or printed replicas that look entirely legitimate in standard capture flows, especially when lighting is flat and resolution is compressed. To a human reviewer or a basic system, everything checks out. But under deeper inspection, these documents betray themselves through missing depth, inconsistent light reflection, and material properties that don’t match genuine substrates. The illusion holds until you look at how the document behaves, not just how it looks.

The takeaway
If you’re only verifying what the document looks like, you’ll miss what it actually is.

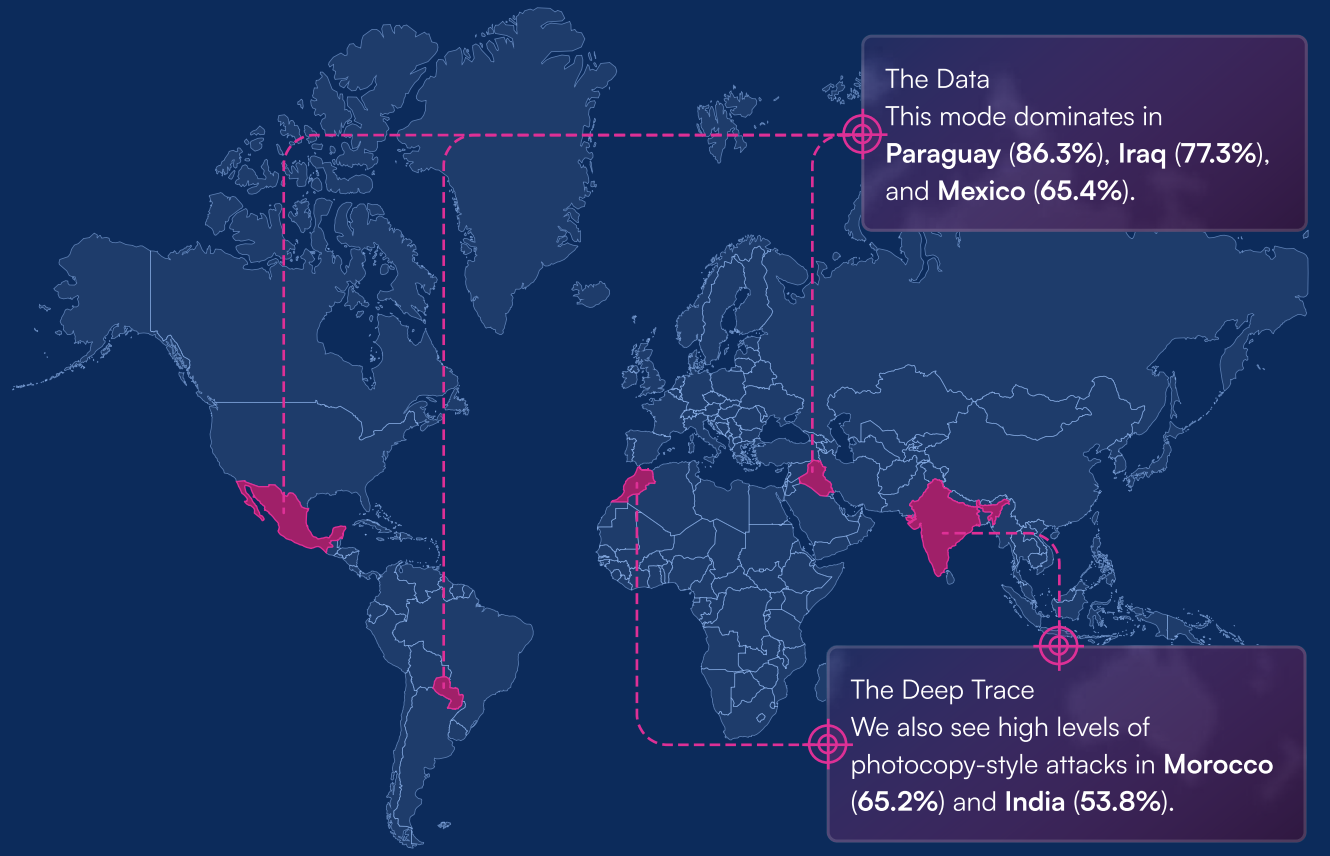

EMERGING THREATS

The AI & Barcode Frontier

What makes this category dangerous isn't volume. It's trajectory. AI-assisted edits are getting cheaper, faster, and more convincing with every iteration, compressing what used to require skill into something that can be scaled. These fakes often pass initial visual inspection because they're designed to look "perfect," but perfection itself becomes a signal. When every pixel is clean, every edge is sharp, and nothing degrades naturally, something is off.

That's where deeper validation starts to matter. Beneath the surface, we consistently find mismatches between what the document shows and what its underlying data claims. Barcodes, in particular, act like a quiet truth serum. When the encoded data doesn't align with the visible information, it exposes a fundamental break in authenticity that surface-level checks miss.

SYNTHETIC & AI-ASSISTED CONTENT

We are monitoring a growing trend of AI-generated document elements. While still a modest percentage of the total, the "quality-to-cost" ratio for fraudsters is plummeting. We treat this as an evolving frontline.

THE BARCODE "TELL"

Particularly in North America, we see documents that look physically flawless but contain **Barcode Anomalies**. Our internal "Verify" solution catches these by cross-referencing the encrypted data strings-a check many automated systems overlook.

The Fraud Attack Adoption Curve

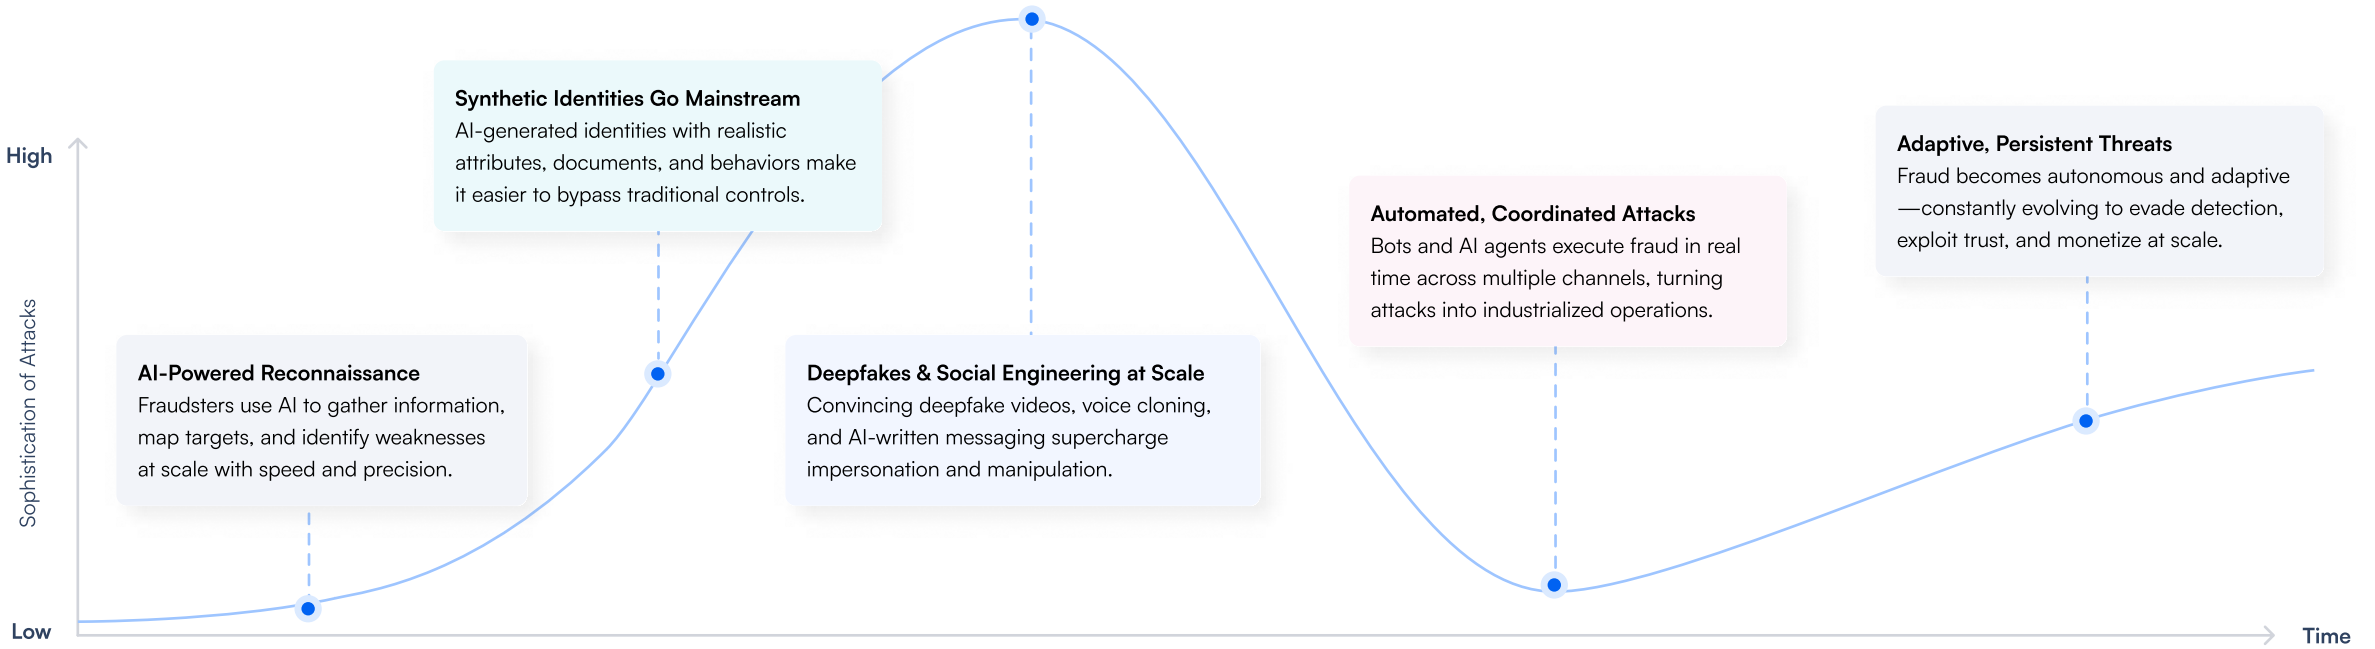

- AI scrapes public data and social media
- Discovers personal and business intel
- Builds targeting profiles and attack playbooks

›

- Generative AI creates realistic fake identities
- Synthetic documents look increasingly real
- Passes basic verification with ease

›

- Deepfake video and voice impersonation
- AI-crafted phishing and fake support
- Emotional manipulation at scale

›

- Bots execute attacks 24/7 across channels
- Credential stuffing, account takeovers, payment fraud
- Toolkits and malware as a service

›

- AI adapts in real time to bypass defenses
- Attacks persist across identities and devices
- Monetization via fraud networks and marketplaces

Key takeaway: Fraud has evolved from isolated incidents to intelligent, dynamic, and scalable operations. Organizations must move from static controls to continuous identity assessment and adaptive defense.

Fraud is no longer an event. It's a continuous system.

The Old Fraud Playbook is Outdated

Fraud Rules Are Breaking

- Device rules (1 device vs multiple) → **outdated**
- Location rules (user must match location) → **broken by modern commerce**
- Velocity/distance rules → **weakened**
- Merchant-of-record shifts (Square example) → **false declines**

Fraud has evolved faster than the rules designed to stop it. All this points to a broader shift underway, namely that fraud is no longer static, isolated, or predictable.It is adaptive. Automated. AI-assisted. Continuously evolving. That means fraud defense must evolve in the same way.

AT A GLANCE

instacart

Case Study: Instacart Modernizes Payment Infrastructure with Microblink

Problem: Manual payment card entry created unnecessary friction during checkout, slowing onboarding and increasing the likelihood of user abandonment.

Solution: Instacart integrated Microblink's payment card scanning technology to enable fast, accurate card capture without manual typing.

Results: The initiative delivered a 5X increase in new users adding payment cards while supporting Instacart's broader payments modernization goals.

From Static Rules to Continuous Intelligence

The rise of AI-driven fraud, shifting digital behaviors, and the growth of always-on commerce are fundamentally changing how organizations think about trust online. Fraud no longer occurs at isolated moments like onboarding or checkout. Instead, it now evolves continuously across sessions, devices, channels, and interactions. Synthetic identities mature over time. Accounts become compromised long after initial verification. AI-generated personas can mimic legitimate users with increasing realism, while automated agents and bots operate around the clock across digital ecosystems that never truly go offline.

This is why identity verification can no longer rely on static, one-time checks alone. Organizations increasingly need continuous identity assessment that evaluates trust dynamically throughout the entire customer journey. Behavioral signals, device intelligence, biometrics, transaction context, and ongoing risk analysis must work together to determine not just who a user claims to be, but whether the interaction remains trustworthy over time.

AT A GLANCE

Case Study: Global Cruise Operator Achieves 90% Increase in Pre-Boarding Verification

Problem: Rising passenger volumes and increasing regulatory requirements were creating bottlenecks in the cruise operator's embarkation and identity verification process.

Solution: The operator integrated Microblink's mobile and on-site document scanning technology into its guest check-in workflows to automate identity verification before passengers reached the port.

Result: The initiative increased pre-boarding verification by 90%, reducing port congestion, manual review workload, and reliance on legacy hardware systems.

The Importance of Know Your Actor

Static verification models built around one-time checks, fixed rules, and isolated signals are increasingly insufficient against adaptive, AI-driven fraud. Organizations now require continuous identity intelligence capable of evaluating trust dynamically over time using layered signals such as biometrics, behavioral analysis, device intelligence, document authenticity, and contextual risk.

At Microblink, we believe the future of identity verification is not simply about determining whether a document or user appears legitimate at a single point in time. It is about continuously assessing whether the interaction itself remains trustworthy throughout the entire customer journey. That's why we believe we are in an era known as Know Your Actor.

In a world of AI-generated identities, autonomous agents, account takeovers, and continuously evolving fraud, the question is no longer simply "Who is the user?" Organizations increasingly need to understand who or what is acting at any given moment, on whose behalf, and with what level of authority. Know Your Actor extends identity verification beyond static onboarding into continuous identity control, helping organizations evaluate trust dynamically across the full customer journey.

AT A GLANCE

Case Study: Tune Talk Modernizes SIM Card Registration with Microblink

Problem: Tune Talk needed to accelerate SIM card registration and customer onboarding while maintaining compliance with Malaysia's telecommunications regulations.

Solution: The company implemented Microblink's automated identity document capture and verification technology to streamline high-volume registration workflows.

Result: Tune Talk reduced SIM registration times by up to 30%, improved agent efficiency, and reduced manual errors across the onboarding process.

Key Takeaways

FRAUD IS BECOMING MORE SPECIALIZED

Fraud no longer behaves uniformly across regions, industries, or attack surfaces. Different markets now experience entirely different attack patterns shaped by local regulations, document ecosystems, and digital behaviors. Organizations should move away from “one-size-fits-all” fraud models and instead build defenses that can adapt regionally and contextually.

STATIC RULES ARE LOSING EFFECTIVENESS

Traditional fraud rules built around location mismatches, device counts, velocity checks, and isolated onboarding events are increasingly unreliable in modern digital ecosystems. Fraudsters continuously adapt to these controls, while legitimate customer behavior has also become more dynamic due to mobile commerce, remote work, and always-on digital services.

TRUST MUST BECOME CONTINUOUS

Identity verification can no longer rely solely on one-time checks at onboarding or login. Trust now needs to be evaluated continuously across the customer lifecycle using layered signals such as biometrics, behavioral analysis, device intelligence, transaction context, and ongoing risk monitoring.

EXPLAINABILITY AND VISIBILITY MATTER

As AI becomes more embedded in fraud prevention systems, organizations need greater transparency into how decisions are made, what attack vectors are being evaluated, and where gaps may exist. Black-box systems without explainability, auditability, or benchmarking create operational and compliance risk.

Why Microblink

Microblink's approach to identity verification is built around a persistent identity layer that connects signals and drives decisions in real time. This enables organizations to detect and stop high-velocity, Gen AI-powered fraud.

Our Identity Intelligence OS captures IDs, biometrics, payment cards, non ID documents and fraud attacks in one flow, delivering decision-ready signals in less than three seconds with no multi-vendor backend handoffs. Every layer is original Microblink IP, from capture and document checks to biometrics, anti-spoofing, and signal correlation. No third-party relays, no black boxes.

KNOW YOUR ACTOR ACROSS THE CUSTOMER LIFECYCLE:

Onboarding → Authentication → Transaction → Monitoring

Glossary

AI-Assisted Fraud - Fraud techniques enhanced or automated using generative AI, machine learning, or synthetic media technologies to create more scalable and convincing attacks.

Barcode Anomaly - A mismatch between the visible information on an identity document and the encrypted or encoded data contained within its barcode, often indicating tampering or forgery.

Behavioral Analysis - The process of evaluating how users interact with systems, devices, or applications to identify suspicious or anomalous behavior patterns over time.

Continuous Identity Assessment - An identity verification approach that continuously evaluates trust across the customer lifecycle rather than relying solely on one-time onboarding or login checks.

Deepfake - AI-generated or AI-manipulated media, including images, video, or audio, designed to realistically imitate a real person or identity.

Device Intelligence - The collection and analysis of device-related signals such as hardware, software, browser, and behavioral attributes used to assess trust and detect fraud.

Document Forgery - The manipulation, alteration, or synthetic creation of identity documents intended to deceive verification systems or human reviewers.

Face Swap - A type of deepfake attack where one person's face is digitally overlaid onto another person in real time or within recorded media.

Injected Stream Attack - A fraud technique where prerecorded or AI-generated video content is digitally injected into a verification session instead of using a live camera feed.

Presentation Attack - An attempt to deceive a biometric or identity verification system using artifacts such as printed photos, screen replays, masks, or manipulated digital content.

Transaction Context - The surrounding information associated with a digital interaction or transaction, including timing, device, location, behavior, and risk signals used to assess trustworthiness.

Velocity Rule - A traditional fraud detection rule designed to flag unusually rapid or repeated activity, such as multiple transactions, logins, or account actions within a short timeframe.