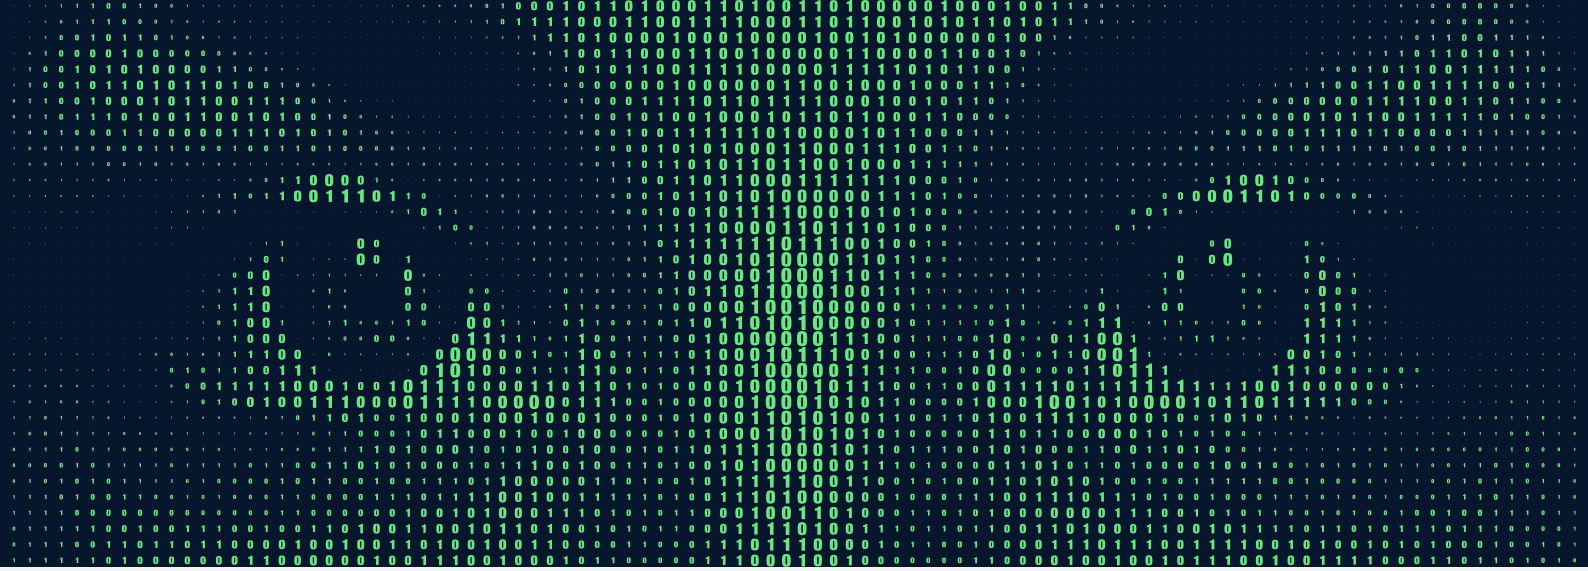
VIDA



VIDA 2026 SEA Digital Identity Fraud Outlook

Designing Identity Defense Against Industrialized Fraud in SEA

www.vida.id



Introduction

Fraud has traditionally been measured by its impact: financial losses, incident counts, and post-event analysis. In today's digital economy, those measures are no longer sufficient. They describe what was lost, but offer limited insight into how fraud is evolving—or why defensive systems increasingly struggle to keep pace.

As Indonesia's economy becomes more digital, identity verification and authentication have become critical infrastructure. These systems now sit at the center of a rapidly changing threat landscape shaped by automation, generative technologies, and the industrialization of scams. Fraud is no longer primarily opportunistic. It is increasingly organized, repeatable, and scaled through tooling that lowers the cost of producing convincing identity artifacts.

This report examines fraud from that boundary.

Rather than reconstructing scams end-to-end, it focuses on what can be observed within verification and authentication flows: how often they are targeted, how attacker behavior changes over time, and how attackers respond as defenses evolve. From this vantage point, fraud appears not as a static problem, but as a dynamic system—one in which attackers learn, adapt, and iterate rapidly.

In this environment, many legacy authentication methods were not designed for the threats they now face. Controls such as SMS-based one-time passwords were built for a world where impersonation required significant effort and identity abuse was largely static. They offer limited resistance to modern attack techniques that combine deepfakes, real-time social engineering, and account takeover workflows—techniques that exploit human trust as much as technical weakness.

As generative fraud accelerates, the challenge is no longer simply stopping individual attacks. Defensive systems must evolve at least as quickly as the tools used to defeat them. Static controls, applied in isolation, age rapidly against adversaries that test, adapt, and return.

The findings in this report point toward the importance of layered defenses—systems that combine multiple signals, adapt continuously, and share intelligence across environments. In a threat landscape defined by speed and adaptation, resilience depends not on any single control, but on how those layers work together.

How to Read This Report

This report is an observational analysis of fraud behavior at the point of identity verification and authentication. It is not a comprehensive accounting of fraud outcomes. Its purpose is to show how fraud behaves at the moment trust is tested—and how that behavior changes as attackers and defenses adapt.

Scope And Vantage Point

The insights in this report are derived from aggregated signals observed within the VIDA verification and authentication network in Indonesia. These signals reflect interaction patterns, artifact characteristics, and system responses during verification and authentication flows.

VIDA does not observe the full scam lifecycle. We do not see social engineering narratives, financial outcomes, or downstream victim impact. As a result, this report does not attempt to quantify losses, reconstruct scams end-to-end, or rank fraud prevalence across regions or industries.

Instead, it focuses on a narrower boundary: the moment where a person, a device, and a system attempt to establish trust.

What The Data Represents

The data presented here is aggregated, anonymized, and analyzed at a behavioral and pattern level. Findings are directional rather than absolute, intended to highlight adaptation and change rather than total fraud volume.

How To Interpret The Insights

The insights in this report describe a dynamic system

- Attackers adapt to defenses
- Defenses reshape attacker behavior
- Generative technologies compress the time between those adaptations

They should be read as documented shifts, not static truths. In a landscape shaped by deepfakes, advanced social engineering, and automation, legacy assumptions about identity and trust are increasingly misaligned with reality.

Chapter 1:
The State of Digital Identity Fraud

Digital fraud is one of the fastest-growing forms of transnational crime. While reported losses continue to rise globally, the more consequential shift lies in how fraud is now produced, distributed, and scaled. Identity—how it is created, represented, and verified—has moved to the center of this transformation.

Global Growth: Scale, Speed, And Professionalization

Global reporting from organizations such as the Global Anti-Scam Alliance indicates that scam exposure now affects a significant share of the global population, with aggregate losses measured in the billions of dollars annually.

More important than headline loss figures is the structural change behind them. Fraud increasingly operates at population scale, enabled by automation and low-cost digital distribution. Scam cycles have compressed, with less time between first contact and exploitation. What once required individual effort now resembles a production system optimized for speed and repeatability.

As a result, modern scams often blend social engineering, technical compromise, and identity abuse into a single flow. Verification and authentication systems have become critical friction points within that flow.

Indonesia: High Exposure, Fast Iteration

Indonesia reflects these global trends in amplified form. Rapid digital adoption, a large mobile-first population, and widespread use of messaging platforms allow scams to spread quickly and adapt rapidly.

On January 10, 2026 - OJK announced that total losses in 2025 reached Rp 9 Trillion, or $534 million with over 125,000 accounts frozen.

Survey-based studies and official reporting suggest:

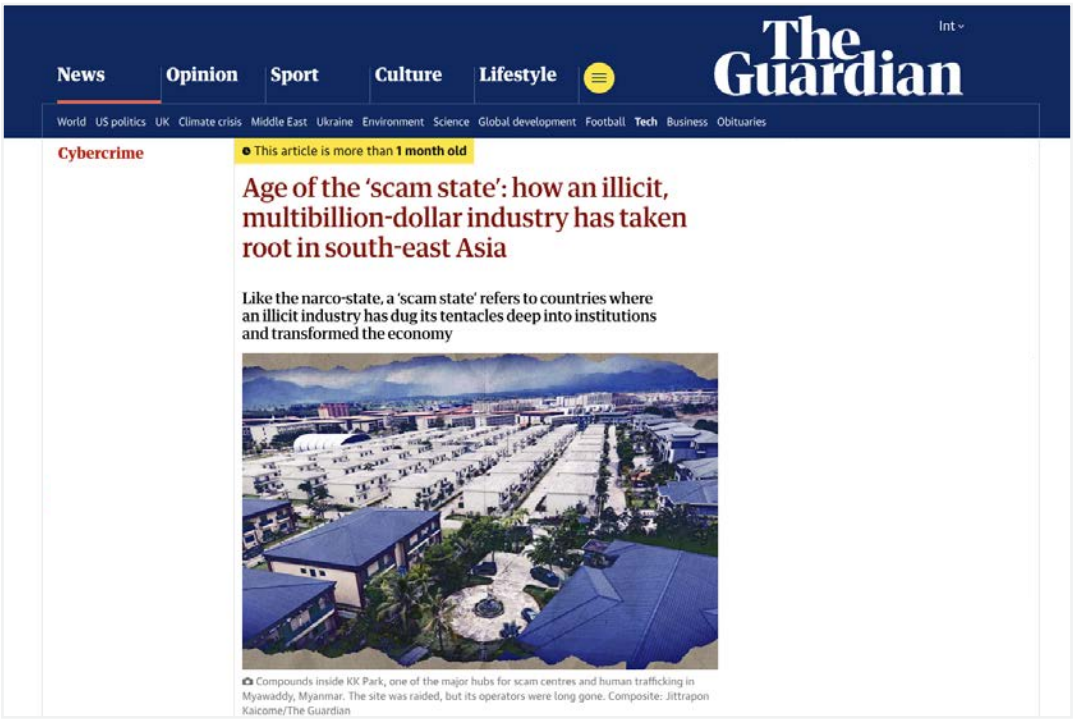
1. According to a 2025 report from the Global Anti Scam Alliance, nearly 66% of Indonesians encounter a scam, facing up to 55 scams a year.

2. A November 2025 report from OJK highlighted the rise of AI related scams including voice cloning and deepfake impersonation between October 2024 and October 2025, amounting to nearly Rp 7.4 Trillion or over $460 million.

At the same time, financial-system reporting captures only a fraction of this activity. Cases that reach banks or regulators represent those that are detected and reported—often after losses occur. A much larger volume of attempts fails or is abandoned at the identity and authentication stage, remaining invisible in traditional statistics.

Organized Scam Operations In Southeast Asia

International reporting has documented the growth of organized scam operations across parts of Southeast Asia, characterized by coordination, shared tooling, and repeatable workflows.



The Guardian

News Opinion Sport Culture Lifestyle

World US politics UK Climate crisis Middle East Ukraine Environment Science Global development Football Tech Business Obituaries

Cybercrime

This article is more than 1 month old

Age of the 'scam state': how an illicit, multibillion-dollar industry has taken root in south-east Asia

Like the narco-state, a 'scam state' refers to countries where an illicit industry has dug its tentacles deep into institutions and transformed the economy

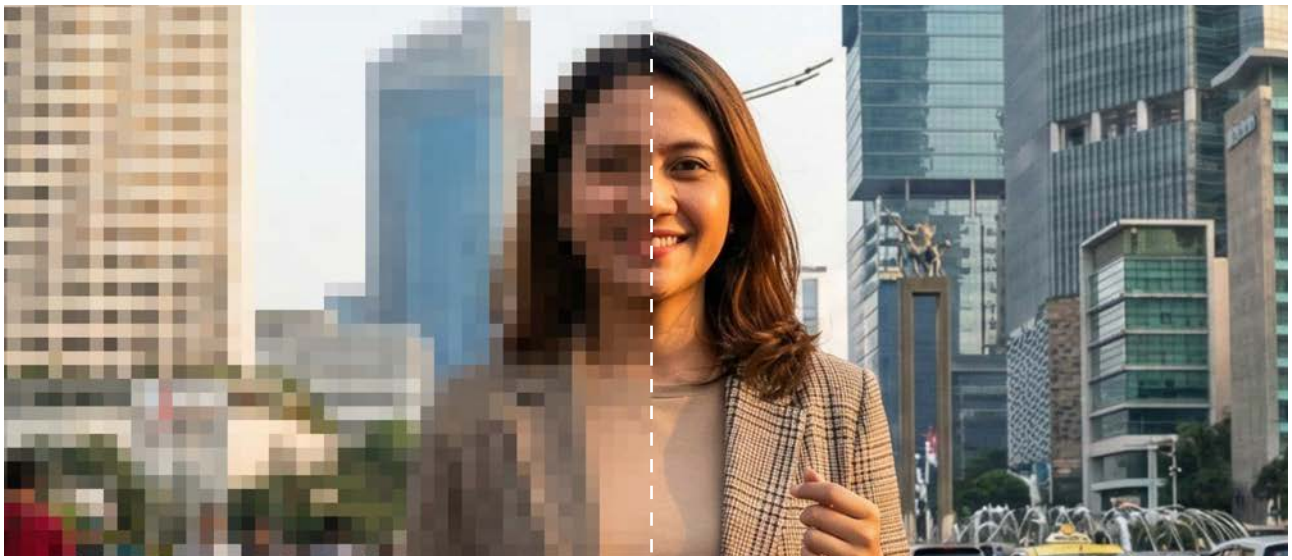
Compounds inside KK Park, one of the major hubs for scam centres and human trafficking in Myawaddy, Myanmar. The site was raided, but its operators were long gone. Composite: Jittrapon Kaicome/The Guardian

These operations illustrate how fraud has evolved beyond isolated abuse into structured systems. Within such environments, identity misuse—through impersonation, synthetic identities, and account takeover techniques—functions as a reusable capability, deployed wherever it increases success.

Why Identity Sits At The Center

Across both global and Indonesian contexts, a consistent pattern emerges: identity is no longer just a gate to be bypassed, but a surface on which attackers test, adapt, and optimize.

As fraud becomes more organized and technologically enabled, pressure on verification and authentication systems increases—not only to block attacks, but to do so continuously against adversaries that learn. Understanding fraud today therefore requires looking beyond reported losses, toward the behavior that unfolds at the moment where trust is tested.



Chapter 2:
The Evolution of Generative Fraud

Fraud attempts observed within the VIDA network have not evolved gradually. They have advanced in stepwise jumps, closely aligned with improvements in generative technologies and the tooling built on top of them.

What appears at the verification and authentication boundary is not simply more fraud, but better fraud—more realistic, more varied, and increasingly difficult to distinguish from legitimate user behavior. These changes are most visible in the quality of synthetic identities, deepfakes, and face manipulation techniques submitted to identity systems.

From Crude Artifacts To Convincing Identities

Earlier fraud attempts relied on low-quality images, obvious manipulation, and repeated reuse of the same assets. They were often easy to detect—not because defenses were advanced, but because the artifacts themselves were poor.

As generative models improved, those characteristics changed quickly. Synthetic faces became sharper and more diverse. Face swaps improved in alignment, lighting, and expression. Artifacts that once looked artificial began to pass casual human inspection and increasingly challenged automated systems.

Within the VIDA network, these shifts appeared as clear inflection points, where the plausibility of fraudulent submissions improved over short periods of time.

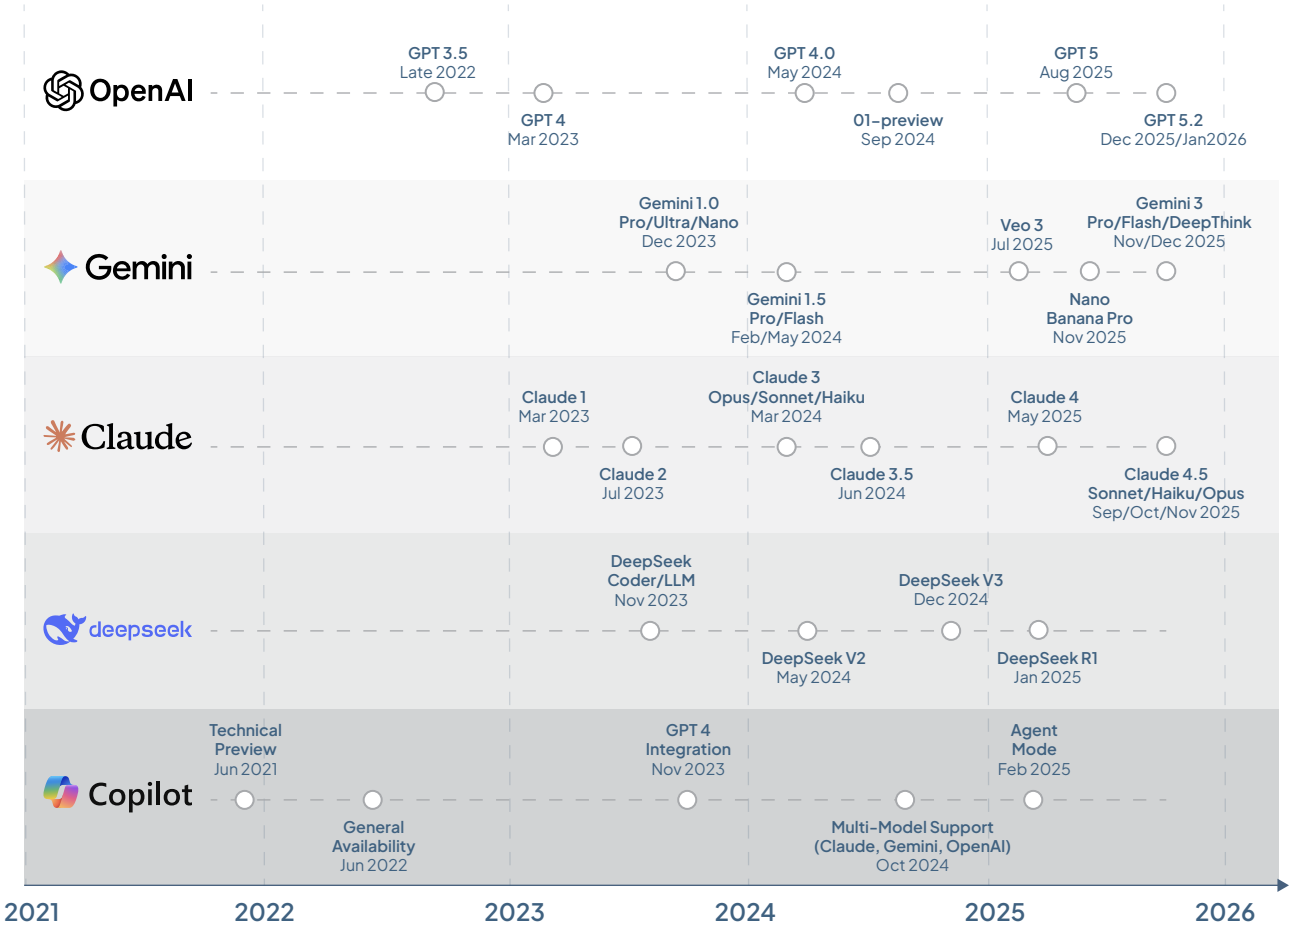
Generative Models As Accelerators, Not Origins

VIDA does not observe which specific tools or models attackers use. What is observable are the effects of successive improvements in generative capability as they propagate through the fraud ecosystem.

The mainstream adoption of AI-generated imagery accelerated with the emergence of accessible image generation platforms such as Midjourney, DALL·E, and the broader Stable Diffusion ecosystem. These tools marked a clear inflection point: photorealistic faces, consistent lighting, and natural facial structure became achievable by non-expert users at scale. Many visual artifacts that once made synthetic identities easy to detect effectively disappeared.

Subsequent advances in large foundation models further accelerated this shift. Improvements in multimodality, context awareness, and automation enabled generative capabilities to be embedded directly into workflows—combining text, image, audio, and video generation into cohesive systems.

AI Model Release Timeline: A Comparative Overview (2021–2026)



As these capabilities improved, downstream tooling adapted rapidly. Face generators, swap engines, and identity assembly pipelines absorbed new techniques with increasing speed. At the same time, cost and accessibility dropped, compressing the time between major model releases and observable changes in fraud attempts.

The result is a compounding effect: each generation raises the baseline of what "convincing" fraud looks like, while shrinking the time attackers need to experiment, fail, and return with improved artifacts.



Chapter 3:
Insight 1: Fraud Quality Evolves in Step with Accelerating Generative Capability

Over the past three years, the quality of fraud observed at the identity verification boundary has not improved incrementally. Instead, it has advanced in distinct step-changes, each redefining what attackers can produce and what defenders can reasonably trust.

Across the VIDA network, these shifts align with rapid advances in generative AI capability. Improvements in image realism, multimodal synthesis, automated post-processing, and accessibility have raised the baseline of fraudulent identity artifacts. As these capabilities became easier to operationalize, attackers adopted them quickly—often faster than defensive assumptions could adjust.

What matters is not any single tool, but the growth of generative capability as a system. As quality thresholds rise and production costs fall, high-fidelity fraud ceases to be exceptional and becomes the default. Techniques that once required manual effort or specialized skill are now increasingly automated, repeatable, and scalable.

This acceleration produces eras rather than trends. Each era reflects a new set of assumptions about what fraud looks like—and which signals remain trustworthy. Earlier techniques do not disappear, but are rapidly overshadowed by more capable approaches.

To understand how fraud has changed—and why static defenses struggle to keep pace—it is necessary to examine these eras directly.

Generative Fraud Eras (2023–2025)

The following pages illustrate how fraud evolved across three distinct eras, using representative, redacted examples observed within the VIDA network.

Era 1: Manipulation & Enhancement (2023)

Definition:
Fraud attempts primarily modify real images using basic manipulation and enhancement to bypass checks.

Era at a glance:

- Low-fidelity face swaps and basic enhancement
- Visible artifacts and repeated templates
- Fraud quality constrained by manual effort
- Visual inspection remained reliable

In 2023, most fraud attempts were limited by source material quality and human effort. Attackers typically began with real photographs and applied rudimentary manipulation to improve pass rates.

Common patterns included rough face swaps with visible blending artifacts, misaligned features, and inconsistent lighting. Basic enhancement, cropping, blurring, and compression were often used to mask inconsistencies. Template reuse was widespread, reflecting limited variation across workflows.

Why this era mattered:
Visual artifacts and inconsistencies remained dependable detection signals.

Representative examples (VIDA network):



Era 2: Photorealistic Deepfakes & Scene Editing (2024)

Definition:
High-quality deepfakes become widespread, expanding fraud from facial manipulation to full scene editing.

Era at a glance:

- Photorealistic deepfakes at scale
- Background and scene manipulation
- Obvious facial artifacts became rare
- Visual realism alone no longer sufficient

By 2024, fraud attempts increasingly featured high-quality deepfakes with minimal visible artifacts. Attackers shifted from faces to full scenes, removing contextual cues that previously supported detection.

Background replacement became common, particularly to avoid recognizable ID or passport photo settings. Full-body alteration and compositing produced images that were internally consistent and visually convincing at a glance.

Why this era mattered:
Detection approaches reliant on obvious artifacts or familiar contexts became unreliable.

Representative examples (VIDA network):



Era 3: Synthetic Identity & Evasion (2025)

Definition:
Attackers move from manipulating real identities to generating synthetic ones—and begin targeting detection systems directly.

Era at a glance:

- Fully synthetic identities
- Inpainting and outpainting for realism with adversarial noise
- Greater variation reduced template detection
- Attacks targeted systems, not just perception

In 2025, attackers increasingly generated fully synthetic identities, combining AI-generated faces with realistic environments and consistent visual details. Inpainting improved facial quality, while outpainting created plausible backgrounds.

Attempts showed greater diversity and less reuse. At the same time, adversarial noise and targeted perturbations emerged, designed to interfere with liveness detection and manipulation classifiers.

Why this era mattered:
Detection systems themselves became the target, compressing response windows and accelerating attacker adaptation.

Representative examples (VIDA network):



Conclusion: What The Eras Reveal

Taken together, these eras illustrate a fundamental change in fraud dynamics. Each era permanently raised the baseline of what attackers could produce, while shrinking the margin for static or assumption-based defenses.

Generative Fraud Eras – Key Differences (2023–2025)

Dimension	**Era 1** **Manipulation & Enhancement (2023)**	**Era 2** **Photorealistic Deepfakes & Scene Editing (2024)**	**Era 3** **Synthetic Identity & Evasion (2025)**
What attackers produced	Modified real images	Highly realistic edited scene	Fully synthetic identities
Source material	Real photos	Real + synthetic composites	Entirely AI-generated
Quality ceiling	Limited by manual effort	High visual realm	High realism + system evasion
Variation at scale	Low (template reuse common)	Moderate	High (low reuse, rapid iteration)
Primary detection signal that failed	Visual artifacts	"Looks real" heuristics	Static models & single-layer controls
Nature of fraud	Opportunistic manipulation	High-fidelity deception	Adaptive, system-aware evasion

As generative capability accelerated, fraud quality improved through rapid absorption of new capabilities into operational workflows. The result is an environment where high-quality fraud is cheaper, faster, and easier to reproduce—and where detection assumptions expire quickly.

This progression explains why defenses based on static rules, isolated checks, or legacy mechanisms struggle to remain effective. Visual realism, contextual familiarity, and even liveness cues have been systematically eroded as reliable indicators of authenticity.

Most importantly, these shifts set the conditions for what follows. When high-fidelity fraud becomes repeatable and inexpensive, it stops behaving like opportunistic abuse and begins to resemble coordinated, scalable operations.

Chapter 4:
Insight 2: Fraud Is Synchronized to Liquidity and Trust Events

Core Finding

Fraud activity within the VIDA network is not evenly distributed over time, nor is it driven purely by technical opportunity. Instead, it clusters around predictable moments of financial liquidity and heightened user trust—when users expect money to arrive and notifications to appear.

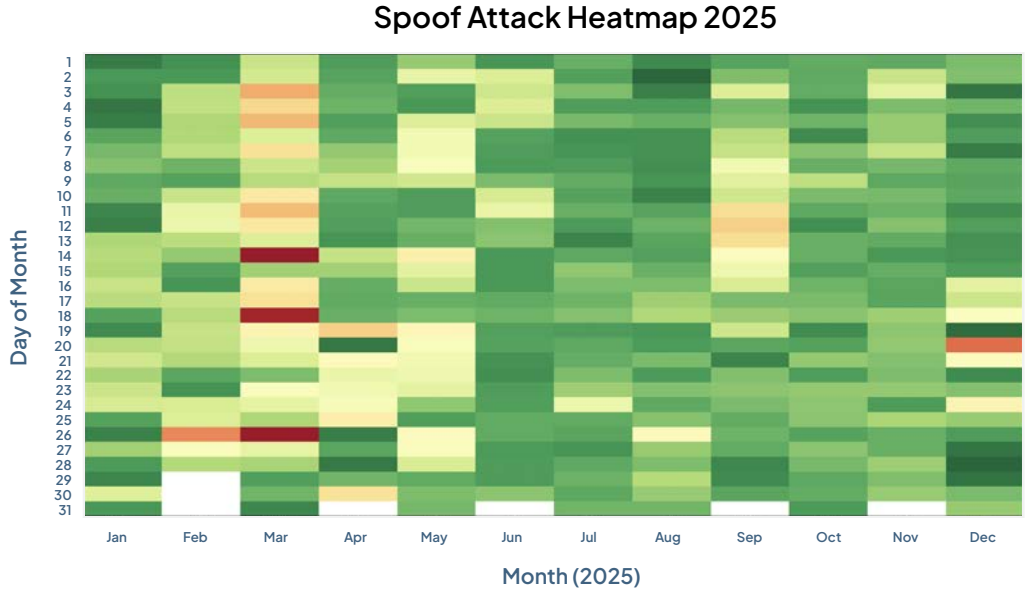
The clearest example in 2025 is the THR (Holiday Allowance) disbursement cycle, which produced the single largest and most sustained concentration of fraud activity observed across the year. This pattern suggests that modern fraud campaigns are economically timed, not reactive.

The THR Super-Spike: Fraud Timed To National Liquidity Events

Aggregated VIDA data shows that March 2025 fraud activity concentrates around specific, institutionally defined milestones in Indonesia's THR process:

- **H-14 to THR** — civil-service THR processing begins
- **Official THR payout dates** — funds released and notifications delivered
- **Late-month private-sector THR deadlines**, overlapping with regular salary cycles

Rather than a single surge, these appear as distinct waves aligned to known payout schedules and user expectations.



Attackers time campaigns for moments when users expect funds, notifications are actively delivered, and banks operate under peak load. Fraud increases not when traffic is merely high, but when liquidity and trust converge.

The “Payday Pulse”: Monthly Liquidity Synchronization

Outside THR, a recurring monthly pattern appears between the 25th and 28th of nearly every month.

Key characteristics:

- Consistent repetition across months
- Mid-week concentration (Tuesday–Thursday)
- Clear deviation from early-month baselines

This aligns closely with standard salary disbursement practices, reinforcing that campaigns are synchronized to **expected liquidity inflows.**

As fraud quality becomes repeatable and inexpensive, **timing becomes the primary multiplier.**

Trust Events As Attack Triggers

Liquidity windows coincide with trusted notifications—bank alerts, employer messages, platform confirmations. Fraud campaigns appear designed to exploit this overlap through impersonation and urgency tied to expected transactions.

Fraud here exploits not just systems, but **human expectation at precisely the right moment.**

Defenses that assume evenly distributed risk will underestimate exposure during predictable surge windows. Static thresholds and uniform authentication fail precisely when risk is highest.

Effective defense requires recognizing **the calendar as an attack surface**—anticipating when risk peaks, not just how attacks appear.

When fraud becomes synchronized to calendars and liquidity, it creates the conditions for **scale and coordination**. Precision timing allows attackers to concentrate effort, reuse assets, and deploy campaigns efficiently across multiple targets.

In this environment, fraud stops behaving like isolated abuse and begins to resemble an **organized production system**—with scheduling, reuse, and optimization as core features.

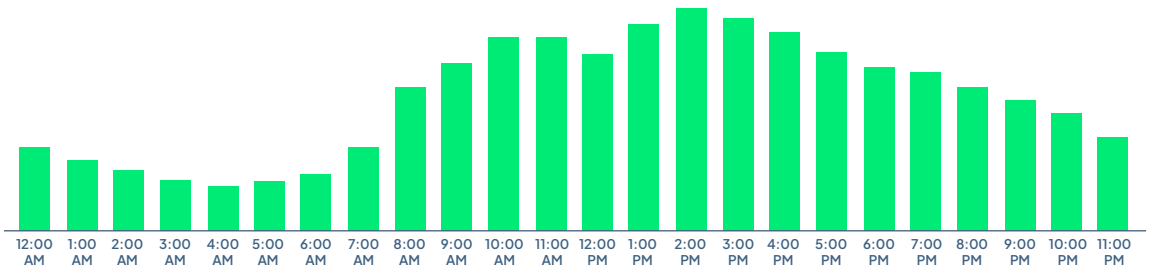
Chapter 5:
Insight 3: Fraud Is Becoming Industrialized

As generative capability raises the quality ceiling of fraudulent identity artifacts, fraud increasingly behaves less like opportunistic abuse and more like a production system. Within the VIDA network, this industrialization shows up in consistent operational patterns—not just higher volumes.

Reason 1: Time-Based Regularity Signals "Work," Not Opportunism

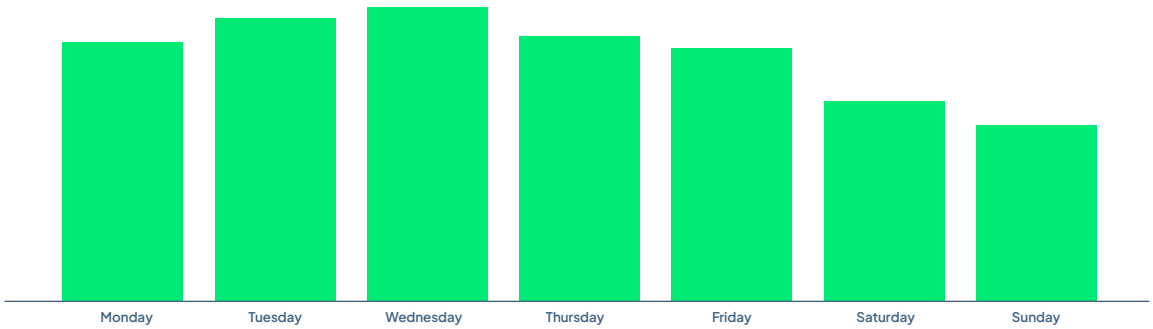
Across the VIDA network, fraud attempts follow **predictable daily rhythms** that resemble operational schedules more than organic user behavior. Activity rises and falls in a way that aligns with standard working hours, with lower activity outside those periods.

Fraud Signals Follow Office Activity Schedules



This is a meaningful shift: rather than being driven purely by chance moments or isolated actors, a portion of fraud traffic appears consistent with organized labor patterns—the kind you would expect from coordinated teams running repeatable processes.

Fraud Signals Follow Weekday Schedule



Implication: detection and response cannot be optimized only for "random" abuse; it must account for **operations that run on schedules.**

Reason 2: Reuse Across Clients Suggests Coordinated Exploitation

A second industrial signal is the reuse of the **same identity artifacts**—the same images, synthetic faces, or identity elements—appearing across different client environments. VIDA has recorded up to 30 instances of the same identity being used in different variations to try to attempt to conduct scam activities.

When identical or near-identical fraud inputs show up in multiple organizations, it suggests more than isolated attempts. It points to **concerted exploitation**, shared playbooks, and re-deployment of assets until defenses change. It also reinforces a network reality: when defenders identify and block a reusable fraud asset, the benefit extends beyond a single organization.

Implication: fraud should be treated as **campaign activity**, and defenses should assume cross-target reuse.

AI-Generated Samples Representative of Trend:

Same Face



Bank A



Bank B



Bank C



Bank D



Bank E



Bank F

Different Backgrounds



P2P A



P2P B



P2P C



P2P D



P2P E



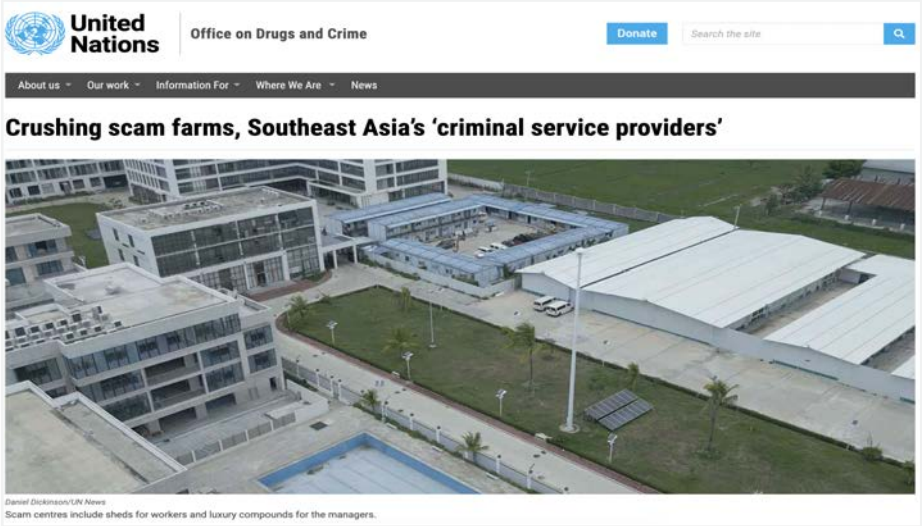
P2P F

Reason 3: Scam "Compounds" In Southeast Asia Reflect Real-World Industrialization

VIDA's signals align with a broader regional reality: Southeast Asia has become a hub for large-scale scam operations that function like industrial facilities—often relying on trafficked labor and structured workflows. UNODC has documented the expansion and global reach of scam centers and the criminal ecosystems around them. Reporting and investigations also describe the scale and persistence of scam compounds across Cambodia and the broader region.

Indonesia is directly exposed to this regional system through repeated cases involving Indonesians trafficked or detained in raids linked to online scam operations in Cambodia.

Implication: the "industrial" patterns observed in verification and authentication traffic are consistent with **industrialized upstream operations**—organized, staffed, and sustained.





Industrialized fraud changes the defensive problem. The goal is no longer just blocking individual attempts, but reducing the efficiency of repeatable workflows—through layered defenses, networked intelligence, and rapid adaptation.

Chapter 6:
Insight 4: Defenses Must Evolve at the Speed of Fraud—or They Will Age Out

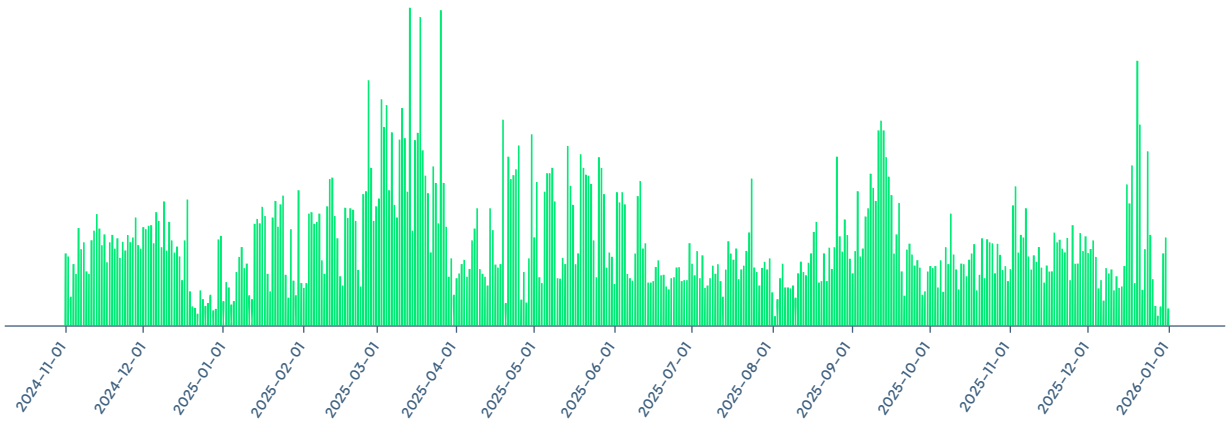
As fraud quality improves and operations become industrialized, the limiting factor is no longer attacker capability—it is **defensive adaptability**. Across the VIDA network, fraud does not disappear when blocked. It shifts form, timing, or technique, often re-emerging once defensive assumptions fall behind.

The defining challenge for security teams today is simple: **static defenses age faster than ever.**

Defensive Gains Are Temporary In A Rapidly Shifting Environment

VIDA consistently observes the same pattern following defensive improvements. After new detection or verification models are deployed, fraud volume and success rates decline. Over time, activity returns—often altered, more evasive, or shifted to different approaches.

This does not indicate defensive failure. It reflects a structural reality: **defenses reshape attacker behavior only while they remain current.** As attackers adapt, fixed advantages erode.



Why Defensive Assumptions Now Decay In Weeks, Not Months

Over the past three years, the pace of state-of-the-art (SOTA) advancement in generative systems has accelerated dramatically. What once took months to meaningfully improve now evolves in weeks. As these cycles compress, attackers gain access to improved realism, automation, and evasion at unprecedented speed.

Why it matters for Fraud: This is your **Defense Window**. In 2023, you had a year to build defenses against GPT 4 class threats. In 2026, you have weeks.

Era	Dominant Model	SOTA Half-Life	Note
2023	GPT 4	**~12 months**	Competitors (Claude 2, Gemini 1) struggled to catch up
2024	Claude 3 Opus/ GPT 4o	**~4 months**	The gap narrowed; competitors caught up in a quarter
2025	Gemini 2/ Claude 3.5	**~6 weeks**	Leadership rotated almost monthly between Google, OpenAI, and Anthropic
2026	Nano Banana/ GPT 5.2	**~2 weeks**	**Current State.** New open-source weights (Llama 4) match closed models almost instantly

At the same time, **the cost of attempting fraud has collapsed**. High-quality synthetic artifacts can now be generated, modified, and retried at negligible marginal cost. Failure no longer carries meaningful penalty; blocked attempts simply become inputs for the next iteration.

Why it matters for Fraud: This quantifies the Volume of Attacks. When cost hits zero, volume hits infinity

Era	Dominant Tooling	Cost per ID Pack	Scale Potential
2022	GPT 3 + Photoshop	**~$15.00**	**Manual:** Required human skill (Photoshop) + API cost.
2023	GPT 4 + Midjourney v5	**~$0.50**	**Semi- Auto:** Good images, but needed manual fixes for text/hands
2024	GPT 4o + Flux.1	**~$0.05**	**Automated:** Flux solved text; scripts could run without human
2026	DeepSeek R1 + Nano Banana	**<$0.001**	**Infinite:** Open-weights allow unlimited local generation for electricity cost only

Definition: The cost (in API credits or hardware) to generate a complete, convincing "Synthetic Identity Pack" (Bio text + Photo + ID Document.

This combination fundamentally changes the economics of fraud. When production costs approach zero and iteration is cheap, scaled fraud becomes not only technically feasible but economically attractive. For organized fraud rings, this environment rewards volume, persistence, and experimentation over precision—allowing attacks to run continuously until success, with predictable returns and limited downside.

Together, these forces reshape the defensive equation. Controls built on fixed assumptions—about visual realism, liveness signals, or user behavior—lose effectiveness rapidly. Legacy mechanisms such as SMS-based one-time passwords, which were never designed for deepfakes, synthetic identities, or scaled social engineering, are particularly vulnerable in this environment.

From Controls To Adaptive Defense Systems

In an environment defined by rapid SOTA turnover and low attack cost, effective defense cannot rely on isolated controls or static rule sets. It must be designed as a layered, evolving system.

Effective defense architectures share three characteristics:

1. **Multiple independent signals rather than single points of failure**
2. **Continuous adaptation as attacker behavior changes**
3. **Network effects, where insights in one context strengthen defenses elsewhere**

The objective is no longer permanent prevention. It is to **increase attacker cost, disrupt repeatable workflows, and reduce the window of exploitability.**

What This Means

Security effectiveness is now temporal. Organizations that treat defense as a fixed capability will see protections age out quickly. Those that design for adaptation—through layered defenses and continuous learning—can remain resilient even as generative fraud systems evolve at speed.

Chapter 7:
Responding to Adaptive Fraud

The evolution of fraud described in this report reflects a fundamental shift in how defenses must be designed and evaluated. As fraud becomes more adaptive, organized, and technologically enabled, static approaches—controls that assume stable attacker behavior—are increasingly misaligned with reality.

Modern fraud does not simply bypass controls; it scales social engineering and accelerates account takeovers, often using synthetic or hybrid identities to blend in. Effective response therefore requires a shift from preventing known attacks to managing a learning system.

Why Static Defenses Fail

Many identity defenses were built for an earlier threat model in which attackers reused the same techniques until blocked. Periodic updates or isolated controls were often sufficient.

That assumption no longer holds.

In environments shaped by automation and generative technologies, attackers can rapidly test defenses, abandon failing approaches, and reappear with modified artifacts, timing, or behavior. Social engineering is increasingly combined with technical abuse to force account access at scale.

Legacy mechanisms such as SMS-based one-time passwords are particularly poorly suited to this environment. Designed for static impersonation, they offer limited resistance to real-time social engineering, SIM abuse, or coordinated account takeover workflows. As a result, regulators and industry bodies are increasingly encouraging stronger alternatives.

Fraud As A Feedback-Driven System

What distinguishes modern fraud is not sophistication alone, but feedback speed. Attackers learn from rejection and friction just as defenders learn from detection.

From the verification boundary, this appears through shortened retry cycles, rapid behavioral shifts, and reuse of successful identity patterns across targets. Synthetic identities accelerate this process by reducing the cost of failure. Every defensive action sends a signal—and every signal is interpreted.

Principles Of Effective Response

More resilient defenses share four characteristics:

Layering over singular checks
No single signal is sufficient. Layered defenses increase the cost of evasion.

Adaptation over static optimization
Controls must evolve with attacker behavior.

Patterns over point detection
Coordinated abuse is revealed through patterns, not isolated events.

Raising attacker cost
The goal is sustained resistance, not momentary blocking.

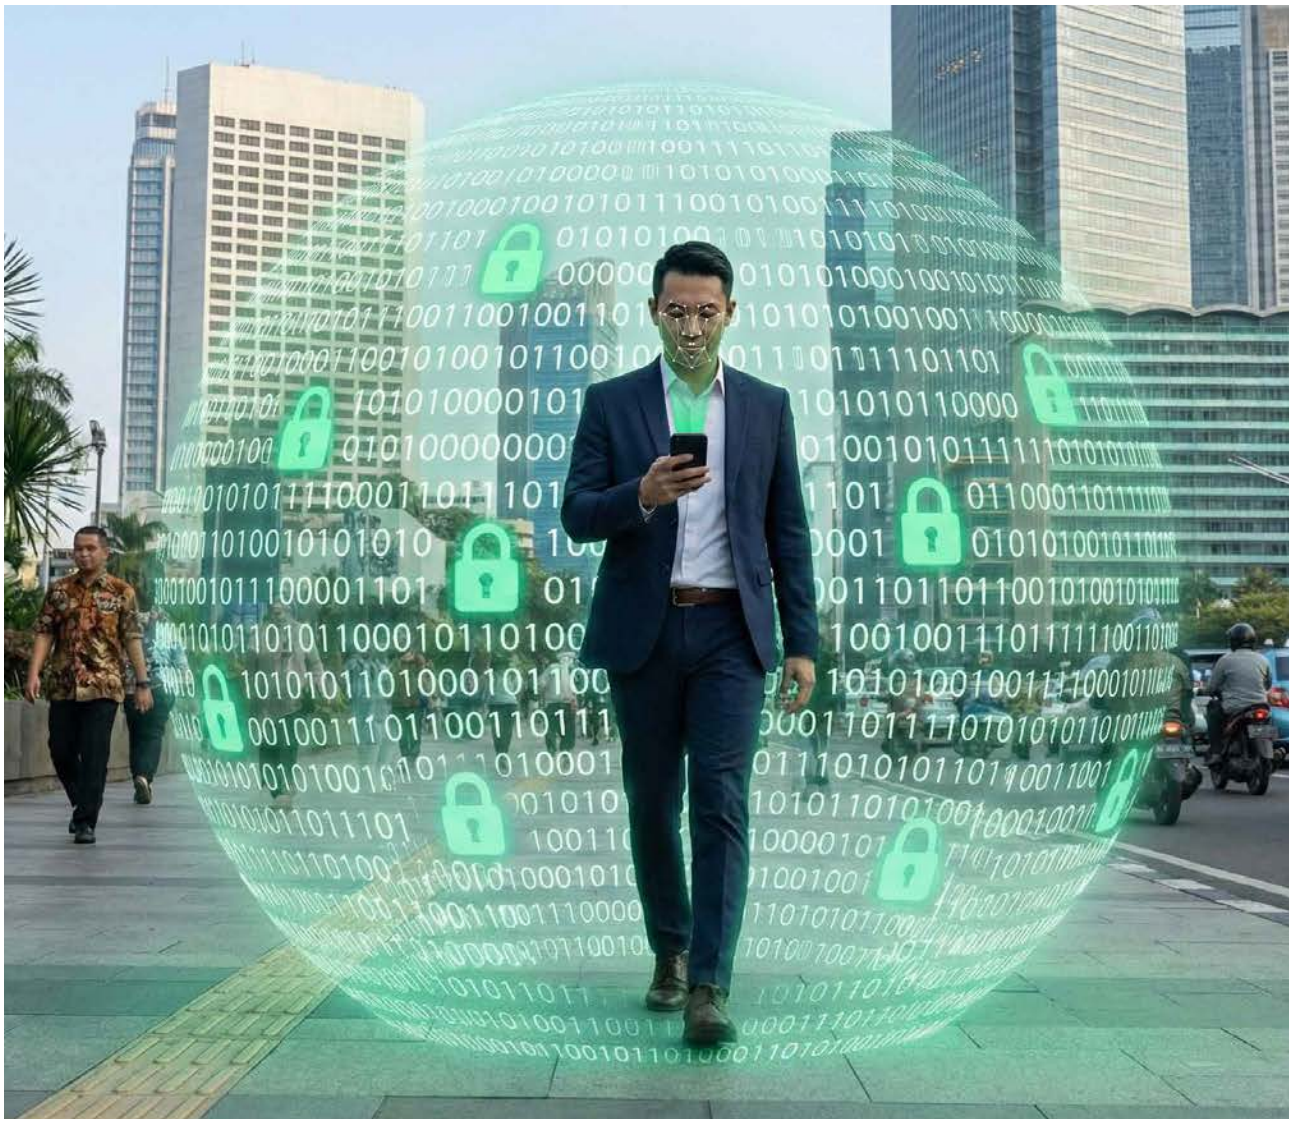
A Networked Perspective On Defense

At scale, techniques that succeed in one environment are reused elsewhere. Networked defenses enable faster recognition of emerging patterns and earlier disruption of coordinated abuse.

Designing For A Moving Baseline

The defining challenge of adaptive fraud is a constantly shifting baseline. What appears anomalous today may be normal tomorrow.

Effective response does not require predicting every new technique. It requires systems designed to assume change and continue functioning as attackers adapt. In this environment, resilience is not the absence of attack—it is the ability to respond faster than attackers can learn.



Chapter 8:
Operationalizing Layers of Defense

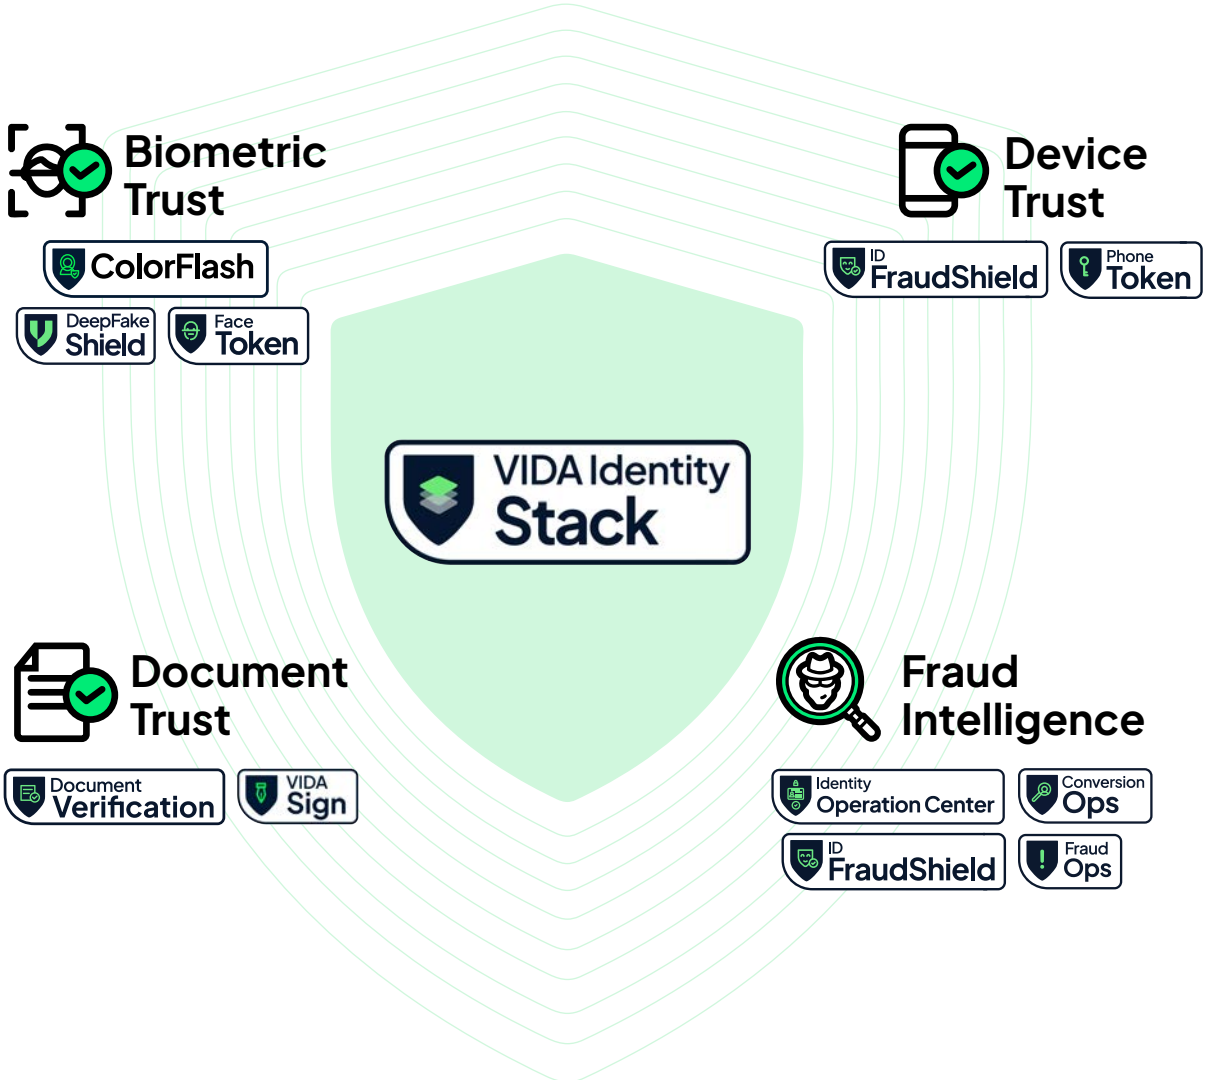
Adaptive fraud systems do not fail at a single point. They probe across multiple surfaces—person, device, document, and behavior—learning where friction is weakest and shifting accordingly. Defending against these systems requires **independent layers of defense that reinforce one another**, rather than isolated controls optimized in isolation.

VIDA operationalizes this principle by unifying multiple verification layers into a single adaptive system—designed to withstand attackers who learn, iterate, and coordinate at scale.

A Layered Response To Multi-Layer Attacks

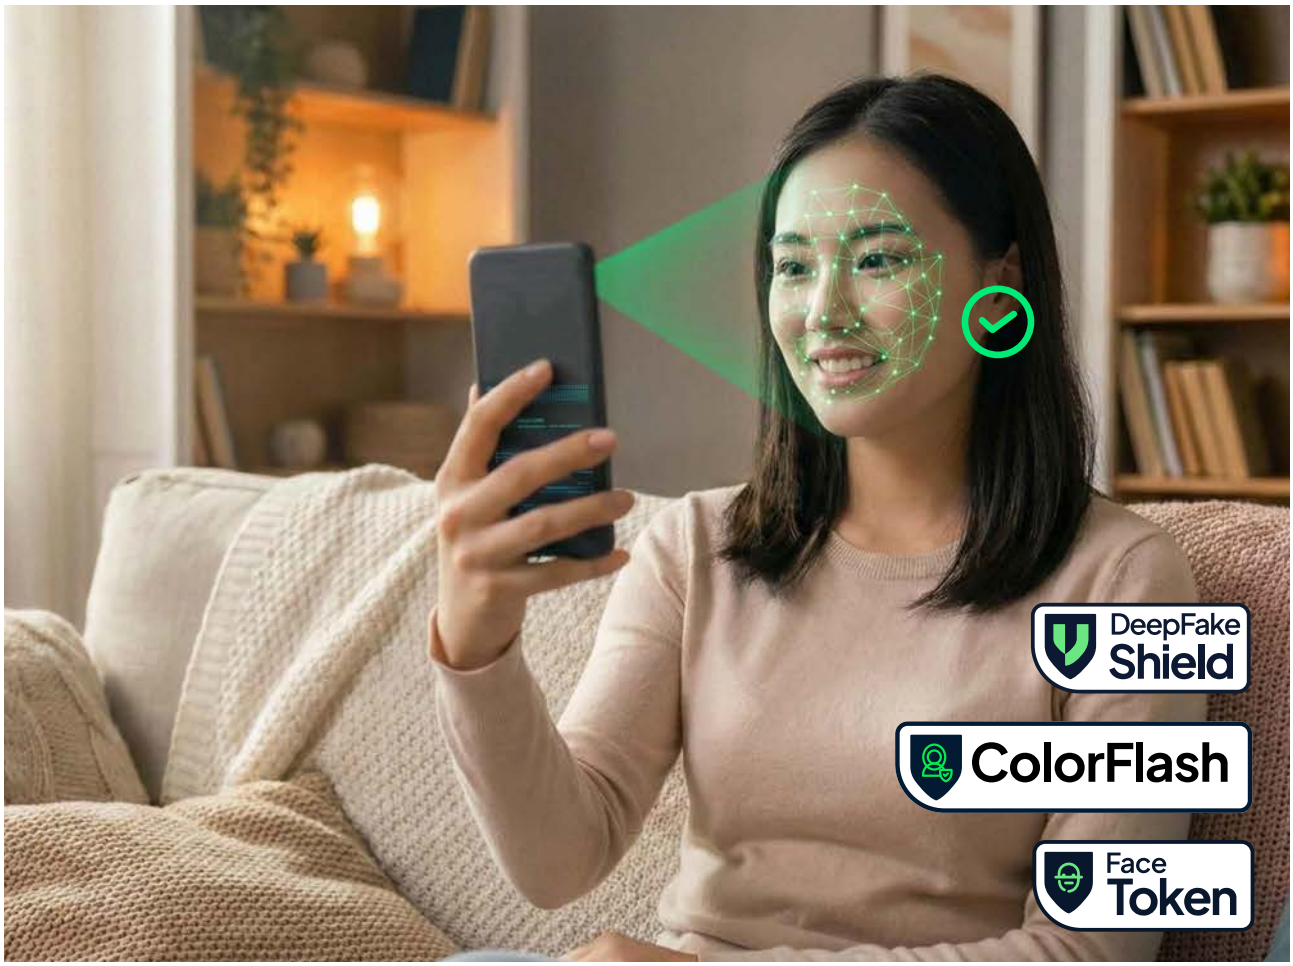
AI-powered scams rarely rely on a single tactic. A synthetic face may appear convincing, but only if the device environment allows it. A legitimate device may still be used to submit manipulated documents. Behavioral signals that appear benign in isolation often reveal coordinated abuse when observed over time.

VIDA's identity stack is structured around four interdependent layers: **Verified Person, Verified Device, Verified Document, and Fraud & Scam Intelligence**. Each answers a distinct trust question. Together, they raise attacker cost, slow iteration, and surface patterns that single checks cannot detect alone.



Layer 1: Verified Person

Is the person real — or an AI-generated proxy?



At the person layer, attackers increasingly rely on face swaps, reenactments, and manipulated selfies to scale social engineering and impersonation.

VIDA addresses this layer through **ColorFlash, FaceToken,** and **DeepFake Shield**, which establish **live human presence** rather than relying on facial resemblance alone. These capabilities detect non-live submissions, AI-generated facial artifacts, and replayed or manipulated biometric inputs.

This layer helps stop fraud that depends on convincing human impersonation—such as fake executive approvals, family-in-distress scams, and call-center-driven deception. As generative realism improves, however, person-level checks must be reinforced by other layers.



Layer 2: Verified Device

Is this interaction coming from a real device — or a controlled environment?

Scalable fraud operations rarely rely on consumer-grade devices. Instead, they use emulator farms, virtual phones, proxy infrastructure, and remote-access tooling that enable rapid reuse and coordination.

VIDA evaluates device trust using **PhoneToken** and **ID FraudShield**, identifying signals associated with emulation, SDK manipulation, injection attacks, and remote-control tools. This layer is particularly effective against automated account takeover workflows and OTP interception schemes—areas where legacy mechanisms such as SMS OTP offer little resistance.

By anchoring interactions to real, uncompromised devices, this layer disrupts the infrastructure that makes large-scale fraud economically viable.

Layer 3: Verified Document

Is the identity document real — or AI-generated?

As generative tools reduce the cost of forgery, attackers increasingly submit synthetic or manipulated documents—often paired with otherwise legitimate biometric or device signals.

VIDA addresses document trust through **Document Verification** and **VIDA Sign**, focusing on authenticating government-issued IDs, detecting tampering, and identifying AI-generated or composited document images. This layer is critical for preventing synthetic identities assembled from mixed real and fake data.

Document verification anchors identity to verifiable sources, ensuring trust does not rest on appearance alone.



Layer 4: Fraud & Scam Intelligence

Is this behavior legitimate — or part of a coordinated attack?

While individual layers assess trust at a moment in time, fraud intelligence evaluates **patterns over time**. This layer connects signals across attempts, users, and organizations to detect coordinated abuse that no single transaction reveals.



VIDA's **Fraud & Scam Intelligence** layer—powered by **ID FraudShield**, the **Identity Operations Center**, **Fraud Ops**, and **Conversion Ops**—analyzes velocity patterns, repeated identity misuse, known scammer device fingerprints, and cross-platform impersonation attempts.

It also incorporates human-in-the-loop review for high-risk or ambiguous cases, ensuring automation is complemented by expert judgment. In adaptive environments, this layer functions as the system's memory—capturing what attackers reuse, refine, and attempt again elsewhere.

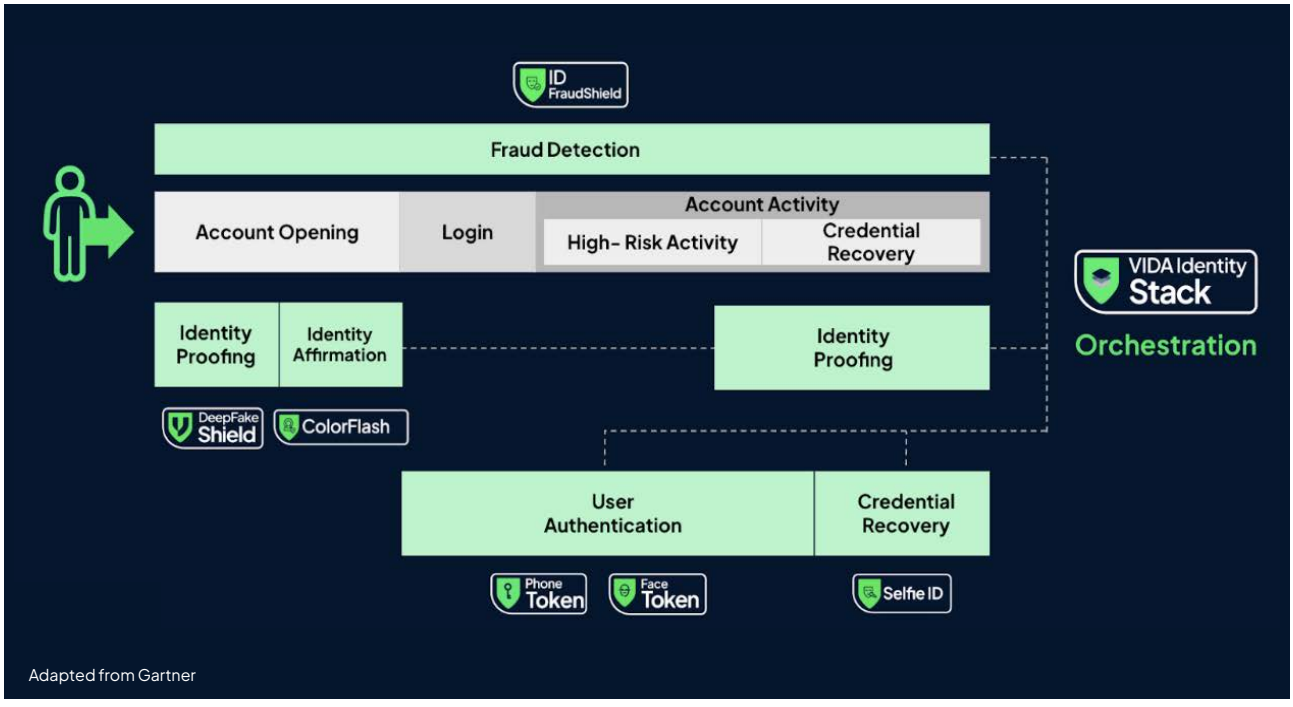
Layered Defense As A System, Not A Checklist

Each layer of defense operates at a different point in the user journey—and challenges a different attacker assumption at that moment. Identity proofing constrains who can enter the system. Authentication governs how access is maintained. Fraud detection monitors how accounts behave once trust is granted. Orchestration connects these layers, deciding when to reinforce trust and when to re-challenge it.

More importantly, these layers reshape attacker economics across the lifecycle. Synthetic identities become more expensive when identity proofing and device trust are enforced at account opening and recovery. Emulator farms lose effectiveness when behavioral signals are correlated across login and high-risk activity. Successful techniques decay faster when shared intelligence shortens the window between discovery and response.

No single control is sufficient at any single stage. Resilience emerges from how defenses are sequenced, reinforced, and re-invoked across the journey—forcing attackers to defeat multiple independent checks while defenders learn from every interaction.

In an environment where fraud systems adapt continuously, layered defense is not redundancy at a single checkpoint. It is coordinated resilience across the entire trust lifecycle.



The implications of this layered approach extend beyond technology choices. They affect how identity risk is governed, measured, and prioritized at the organizational level. For security leaders, the challenge is no longer selecting individual controls, but designing systems that remain effective as attackers adapt.

Chapter 9:
What This Means for Security Leaders

For security leaders, the implications of adaptive fraud extend well beyond individual controls or isolated incidents. The challenge is no longer defending against a fixed set of techniques, but operating in an environment where attackers learn quickly, coordinate at scale, and time their activity precisely. Identity risk must now be governed as an evolving system.



Identity Is A Primary Attack Surface

Identity systems are no longer peripheral safeguards; they are core infrastructure. As scams scale through social engineering, automation, and synthetic identities, verification and authentication flows become the primary surfaces where attackers probe for weaknesses and test assumptions.

Crucially, this activity is not random. Fraud increasingly concentrates around predictable moments of liquidity and trust—salary cycles, benefit disbursements, credential recovery, and high-risk account actions. Identity systems are targeted not only for what they protect, but for when they are invoked.

For CISOs, identity can no longer be treated as a compliance layer or downstream dependency. It must be managed as a dynamic attack surface, monitored continuously and reinforced selectively as risk shifts across the user journey.

Legacy Controls Are Increasingly Misaligned

Many widely deployed authentication mechanisms were designed for a different threat model. Controls such as SMS-based one-time passwords assume static impersonation and limited coordination. In a world of real-time social engineering, SIM abuse, and accelerated account takeover workflows, those assumptions no longer hold.

Some controls still add friction but no longer provide meaningful assurance—particularly during high-risk moments such as credential recovery or liquidity-driven activity. Regulators and industry bodies are increasingly acknowledging this gap by encouraging stronger, phishing-resistant and context-aware alternatives.

The strategic question is no longer whether to move beyond legacy methods, but how quickly and deliberately to do so without creating new risk concentrations.

Measure Systems, Not Just Incidents

Traditional fraud metrics focus on outcomes—losses, confirmed cases, post-incident analysis. While necessary, these measures lag the behavior that matters most.

Leading indicators appear earlier at the identity boundary, including:

- Changes in retry and failure patterns
- Increased variation in identity artifacts or behavior
- Shifts in timing aligned to known trust or liquidity events
- Rising use of synthetic or hybrid identities

Security teams must be able to observe and act on these signals in near real time. Measuring system behavior—not just losses—provides earlier warning and greater leverage against coordinated abuse.

Design Defenses For Learning Adversaries

Resilient defenses assume attackers will study them, test them repeatedly, and share successful techniques. This has clear implications:

1. **Layering over single controls**
   No individual signal remains reliable on its own.

2. **Adaptation over static thresholds**
   Defenses must evolve as behavior changes.

3. **Patterns over point events**
   Campaigns, reuse, and timing matter more than individual attempts.

4. **Cost over perfection**
   The goal is sustained resistance, not momentary blocking.

This requires moving from point solutions to orchestrated identity systems—where controls are sequenced and reinforced dynamically across account opening, login, high-risk activity, and recovery.

Leverage Scale And Shared Intelligence

Adaptive attackers benefit from coordination; defenders operating in isolation face structural disadvantages. Techniques that succeed in one environment are quickly reused elsewhere.

Where possible, security leaders should favor approaches that enable shared learning across journeys and organizations. Faster recognition of emerging patterns shortens exposure windows and reduces the effectiveness of industrialized fraud operations.

Reframing Resilience

In an adaptive fraud environment, resilience is not the absence of attack. It is the ability to absorb pressure, adapt defenses, and respond faster than attackers can iterate.

For security leaders, this reframing is essential. It shifts the objective from eliminating fraud entirely to managing identity risk in a way that remains effective as attackers coordinate, optimize, and evolve.

The organizations best positioned to succeed will be those that design identity defenses for change, timing, and orchestration—rather than assuming stability in a system defined by constant adaptation.

Chapter 10:
Final Thoughts: Designing for a Moving Target

The evidence across this report points to a clear conclusion: digital fraud is no longer episodic, opportunistic, or evenly distributed. It is adaptive, coordinated, and increasingly synchronized to moments of trust, liquidity, and operational strain.

Generative technologies have raised the baseline of fraud quality. Calendar-based timing has amplified its impact. Industrialized operations have made it repeatable and scalable. In this environment, defenses built on static assumptions—about identity, behavior, or attacker capability—will continue to age out.

Effective identity defense now depends on orchestration rather than isolated controls. It requires systems that can apply the right layers of verification at the right moments, adapt as attacker behavior shifts, and learn across attempts, journeys, and organizations. The goal is not permanent prevention, but sustained resistance—raising attacker cost, shortening exposure windows, and disrupting repeatable workflows.

For security leaders, the mandate is clear. Identity must be treated as a living system, not a fixed gate. Risk must be anticipated, not merely detected. And resilience must be designed for change, not stability.

In a threat landscape defined by speed, coordination, and adaptation, the organizations that endure will not be those with the most controls—but those with defenses that evolve as quickly as the fraud they face.



Indonesia | Singapore | India | Malaysia