esentry

EYE OF HORUS

Protecting What Matters. Illuminating What Hides.

H1 Report 2026

Leadership Pulse

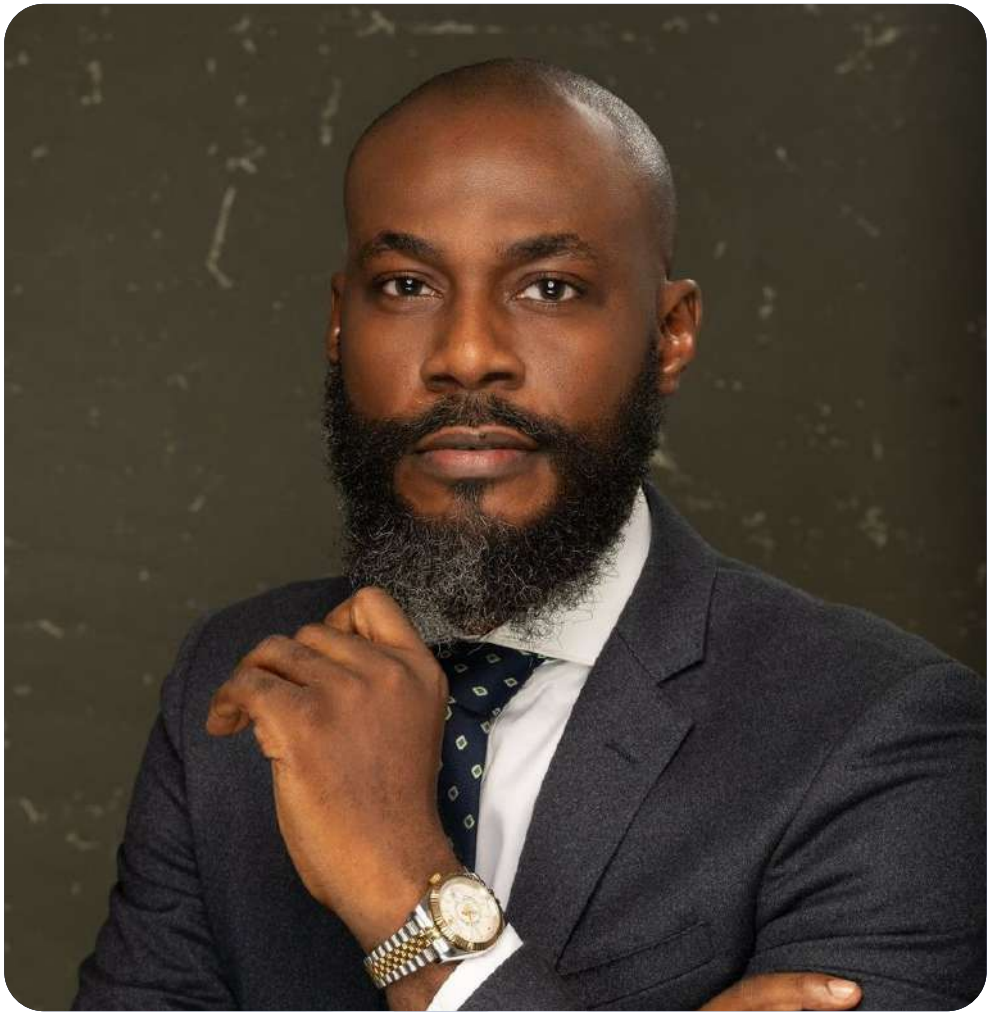

The **Eye of Horus** was many things to the ancient Egyptians: a symbol of **royal protection, a ward against chaos,** and above all, the promise that **nothing, no threat, no act of darkness,** could escape the watchful gaze of those tasked with keeping order.

We chose this name deliberately. The first half of 2026 tested that promise directly, here at home. Nigeria found itself firmly in the spotlight this half year, drawing attention from both highly sophisticated threat groups and opportunistic script kiddies alike. Banks, government platforms, regulatory bodies, and financial institutions all felt the pressure. Weak points that had long been overlooked, unpatched systems, over-permissioned access, and gaps in basic monitoring were the ones that got found and exploited. Attack surfaces widened with every new cloud workload, and the gap between detection and impact kept shrinking. This report is our answer to that environment.

What you read today reflects the combined work of six elite teams, each trained to view the digital world differently, forming a full-spectrum view of every threat, every intrusion, every risk that touched our clients' environments this half year. It is grounded in what we saw in our clients' systems and the global industry, not in talking points.

We also went outside our own walls this time, incorporating perspectives from a respected name in global cybersecurity threat intelligence. Their read on the macro landscape backs up what we are seeing on the ground: attackers are moving faster, automating more of the kill chain, and targeting the seams between cloud and on-premises systems with growing precision.

Their perspective doesn't just add color to this report; it confirms that what we have documented in Nigeria is part of a pattern playing out globally, which makes the case for vigilance even harder to ignore.

> "Read carefully. The Eye sees everything."
>
> **Gbolabo Awelewa**
> Chief Business Officer, esentry

Executive Insight

Africa's rush toward **AI** is outpacing its ability to secure what it builds. Organizations deploying self-correcting agents, automating faster, and cutting headcount are quietly expanding an attack surface most cannot yet see, let alone audit. The speed is real, and so is the exposure. The second half of this year and the year after will make that gap impossible to ignore. We did not wait for that gap to widen before acting on it. Fraud monitoring engagements grew meaningfully this year, and the work that mattered wasn't the alert; it was tracing each case to root cause so it couldn't recur.

That same instinct to see a gap before it becomes a crisis is what led us to look closely at how Africa talks about cybersecurity in the first place.

Too much of that conversation is built for enterprises that already have budgets and boards asking questions. **Startups** and **SMEs** are just as exposed without either. Part of our role as the **Eye of Horus** is noticing what the rest of the market overlooks. Our Startup Bouquet exists for exactly that reason, security priced to fit where a business is today, not once it's large enough to be a target.

Our expansion into **Ghana** reflects the same conviction at a regional scale.

A market tightening its regulatory posture needs partners who build resilience underneath compliance, not vendors who simply check boxes. And because we do not believe in warning the market about a risk we haven't solved ourselves, we launched our own agentic cybersecurity capability this year, built with human sign-off at every high-impact decision point.

This is a record of a company trusted with more because it earned that trust the year before, and the year before that. That's the only growth strategy we have ever had.

Ebun Williams
Chief Operating Officer, esentry

Table of Contents

H1 Threat Landscape

4,700+

WEEKLY ATTACKS

Nigeria national average, H1 2026

281,500

BREACHED ACCOUNTS

Nigeria, Q1 2026 (34th globally)

35+

ORGS BREACHED

ByteToBreach, 26+ countries

1,369

FIRMS UNDER NDPC PROBE

Data Privacy Violations, 2026

Executive Summary

The cyber threat landscape during the first half of 2026 continued to evolve in sophistication, scale, and operational impact. Threat actors increasingly shifted from isolated attacks toward coordinated campaigns leveraging software supply chains, artificial intelligence, automation, credential marketplaces, and social engineering to maximize reach and profitability.

Ransomware remained the most disruptive cyber threat globally, while financially motivated cybercrime groups demonstrated increasing capability to weaponize trusted software ecosystems. Simultaneously, the rapid adoption of Artificial Intelligence introduced new attack surfaces as adversaries exploited AI services, AI credentials, autonomous agents, and AI-assisted malware development.

Africa experienced a significant increase in cyber activity, recording nearly 3,000 attacks per organization every week, highlighting the continent's growing attractiveness to financially motivated and opportunistic threat actors. Within Nigeria, cyber adversaries continued targeting government agencies, financial institutions, technology companies, telecommunications providers, oil and gas organizations, and public-facing digital services, with government institutions remaining the most targeted sector.

Key Highlights for H1 2026

- **Ransomware** has established a new, permanently elevated operational baseline rather than returning to historical volumes; operators increasingly combine double extortion, pre-encryption data theft, supply-chain delivery, and destruction of forensic evidence.
- **Software supply-chain compromise** (exemplified by the TeamPCP campaign against Trivy, LiteLLM, Bitwarden CLI, Telnyx SDK, Checkmarx KICS, and numerous npm packages) means organisations increasingly inherit risk through trusted dependencies rather than direct compromise.
- **AI** is now foundational criminal infrastructure: our assessment found a 640-fold increase in demand for premium AI accounts on criminal marketplaces between Q4 2025 and Q1 2026.
- **Credential theft**, specifically verified, MFA-enabled, and KYC-bundled identity packages, has overtaken traditional malware as the preferred initial-access technique, and synthetic identity fraud is expanding the problem further by removing the victim needed to trigger detection.
- **ByteToBreach** is a HIGH-threat, financially motivated actor active since June 2025, claiming breaches across 35+ organisations in 26+ countries.
- **Nigeria's active threat actor** list (XForum Bot, AlQe3Qa3, DitzzXploit, XYZEAZ, Incransom, Killsec) continues sustained targeting of government, financial services, technology, and oil & gas
- A live campaign abusing genuine Trello board-invitation infrastructure to deliver credential-harvesting links confirms this assessment's own H2 2026 outlook on trusted-SaaS-platform abuse is already materialising, ahead of schedule.

Ghana Threat Landscape

Ghana's digital transformation continued to accelerate through H1 2026, driving growth in digital financial services, cloud adoption, e-government platforms and AI-enabled tools. This expansion has widened the country's attack surface, and esentry's market expansion into the Ghanaian market in 2026 provides an opportunity to examine how this threat environment compares to patterns already observed across esentry's existing African operations.

Ghana's national Computer Emergency Response Team recorded more than **3,500** confirmed cybersecurity incidents in Q1 2026 alone, with increasing reports of malware activity, ransomware attempts and attacks against critical information infrastructure. It has also been documented that **352 cyber fraud** cases linked to fraudulent online investment schemes occurred between January and **June 2026,** with victims losing a combined **GH¢3,429,447** in that period.

The threat landscape remains largely financially motivated, with **credential theft, phishing, business email compromise (BEC), mobile money fraud,** and attacks against **internet-facing systems** dominating reported activity.

Key Targeted Sectors

Financial Services
the primary target, driven by high transaction volumes and sensitive customer data.

Government Institutions
targeted through phishing and credential harvesting aimed at publicly accessible systems.

Telecommunications Providers
targeted for subscriber data, mobilemoney platforms and account takeover opportunities.

Educational Institutions
exposed due to the volume of personally identifiable information held.

Critical Information Infrastructure (CII) Operators
under sustained pressure given their role in essential national services.

Prominent Threats Observed

Phishing campaigns
targeting banking customers, corporate employees, and government personnel.

Mobile Money fraud
including SIM swap attacks and fraudulent payment requests.

Credential Theft
via fake login portals and social engineering.

Web Application Attacks
exploiting internet-facing services.

Ransomware Activity
targeting organizations providing essential services.

AI-assisted Social Engineering
increasing the realism of phishing and impersonation attempts.

Threat Outlook

A Regional Lens: How Ghana Compares

The sector-targeting pattern reported in **Ghana - financial services, government, telecommunications,and critical infrastructure**, closely mirrors what **esentry's SOC and DFIR teams** have observed across its established Nigerian operations over the same period. This suggests these dynamics are regional rather than country-specific, shaped by shared factors: **rapid mobile money adoption, expanding e-government services, and uneven security maturity** across mid-sized enterprises.

Operationally, this **consistency means detection use cases**, **threat-hunting playbooks**, and fraud-pattern intelligence developed against Nigerian telemetry are likely to transfer directly to the Ghanaian market, rather than requiring development from zero. As esentry's operational footprint in Ghana matures, future reports will incorporate esentry-derived findings, incident data, and case studies specific to this market.

Threat Outlook

Ghana's threat landscape is expected to remain driven by financially motivated cybercrime through the remainder of 2026. As digital services and cloud adoption continue to expand, **identity-based attacks**, **phishing, ransomware**, **BEC**, and **payment-ecosystem fraud** are likely to increase in frequency and sophistication. **Financial services**, **government, telecommunications**, **education**, and **CII operators** are expected to remain the most targeted sectors.


Security
Operations
The All-seeing Eye

Security Operations – The All-seeing Eye

3.5M+
Total alerts processed

11 min
Median detection time

5,065
EDR threats blocked

The Monitoring & Alerting team is the first line of sight, the literal eye in the **Eye of Horus framework**. Operating 24 hours a day, **365 days a year**, our analysts maintained continuous coverage across all client environments, processing a record of over **3.5M security alerts** during H1 2026.

A Look Into H1

H1 2026 was the most operationally intensive monitoring period esentry recorded. The **SOC** processed over **3.5 million security alerts** across all monitored client environments, spanning **web application attacks**, **endpoint compromise attempts**, **network intrusion activity**, and **automated attack infrastructure**. Median time from detection to escalation for confirmed threat activity was **11 minutes.**

The defining characteristic of the H1 threat environment was not sophistication in isolation; it was scale combined with patience. Adversaries ran continuous, automated pressure across every exposed surface simultaneously: probing for the one configuration that had not been patched, the one credential that had been reused, the one service left publicly accessible without restriction.

The question they were asking was not "can I break through this wall," but "which wall has a door that was left open?"

Threat Volume

The first half of 2026 showed that organizations are increasingly being challenged at their internet-facing perimeter.

Attackers are leveraging automation to continuously probe exposed services, searching for vulnerabilities, weak credentials, and misconfigurations that can provide an entry point. As a result, web applications remain the primary target, generating significantly more malicious activity than endpoint-based attacks.

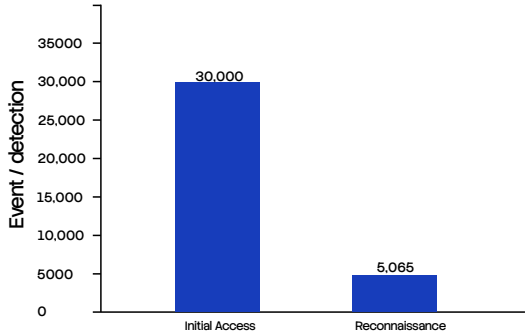

Key detection and blocking volumes across monitored environments, H1 2026.

Attack Types: What Is Showing Up and Why It Matters

Web Application Attacks: 30,000 Events

Web application attacks were the dominant event category in H1. **SQL injection** and **XSS** lead because they are automated and cost almost nothing to run at scale. The category carrying the highest operational risk was **perimeter device exploitation** — specifically **FortiGate SSL-VPN and Ivanti Connect Secure infrastructure** targeted against critical **CVEs** disclosed in Q1 2026. The pattern was deliberate targeting of specific appliance fingerprints, not opportunistic scanning, with exploitation attempts appearing in monitored environments within **72 hours** of public disclosure.

Endpoint:5,065 Threats Blocked

The **endpoint detection layer blocked 5,065** threats across monitored environments. Commodity malware and hacktools comprised the largest share of opportunistic **infostealers**, **RATs**, and **dropper frameworks.** The more operationally significant category was stealth and persistence tooling: **640 detections** of techniques designed to evade signature-based detection and blend with normal system activity, representing cases where the attacker was establishing a foothold rather than deploying off-the-shelf malware.

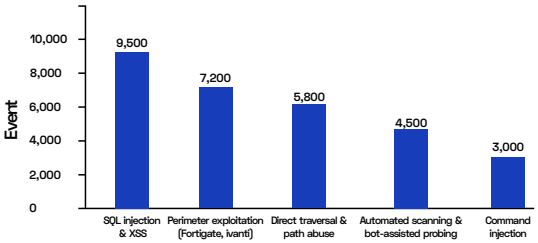

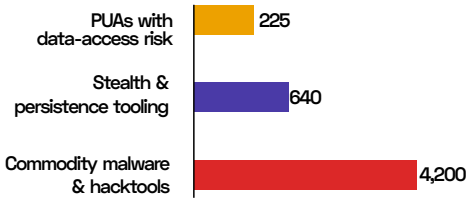

EDR blocked threat breakdown by category, H1 2026. Total:5,065 detections.

Automated Bot Traffic — 5,000 Blocked

5,000 automated bot attack attempts were blocked across monitored environments in H1. Dominant patterns: credential stuffing against **authentication endpoints, account enumeration via API abuse, and scraping campaigns** targeting financial data. Bot infrastructure has matured — payloads rotate across residential IP ranges to avoid reputation-based blocking and operate on timing profiles designed to blend with legitimate peak-hour traffic. Rate limiting and IP reputation controls alone are insufficient countermeasures against this generation of tooling.

The AI Dimension

AI is now a material operational factor on both sides of the detection problem.

On the adversary side: **phishing** and **social engineering** lures intercepted in H1 showed contextual precision absent in prior periods, messages referencing real business conventions and payment cycles, with accuracy consistent with AI-assisted content generation. The barrier to crafting a convincing, targeted campaign has effectively collapsed.

On the defensive side: **AI-assisted triage**, **automated alert correlation**, and **behavioral anomaly detection** make it operationally viable to process **3.5 million alerts** at the quality level required to surface the threats that matter. The human analyst role has shifted toward investigation and judgement, the irreplaceable layer. The two work together; otherwise, neither works at the required scale.

Security Operations – The All-seeing Eye

From Playbooks to Reasoning

That shift deserves its own explanation because it is not the same thing security teams have called automation for the last decade.

SOAR automates known workflows. If an analyst had already written a playbook for a scenario, **SOAR** could execute it. It could not reason about a novel alert or adapt to a pattern nobody had scripted for. **Agentic AI** is a different category of system. Rather than executing a fixed sequence of steps, an agent plans a **multi-step investigation**, pulls context from across the **security stack-endpoint**, **identity**, **cloud**, **email**, **network**, forms a judgement, and takes or recommends action within a defined scope of autonomy.

The industry data on this is now substantial enough to take seriously rather than treat as a vendor promise. **Microsoft's security team** reports agents automating **75%** of **phishing** and **malware investigations internally** and completing vulnerability exposure assessments that previously took a full day of engineering effort in under an hour.

Cisco recorded comparable results at **Cisco Live Amsterdam** in February 2026: an **agentic triage framework** processed **179 incidents** generated from a heavily unmanaged conference network, correctly dismissed **176** as false positives, and surfaced three genuine threats for analyst review, work that would otherwise have consumed most of an analyst's week.

None of this is a one-sided improvement, and we will not present it that way. The same period that made **agentic SOC** capability commercially viable also made AI-assisted attacks materially easier to run.

SoSafe's 2025 Cybercrime Trends report found **87% of security professionals** worldwide had already encountered an AI-driven attack within the prior year, **91%** expected a significant increase going forward, and only **26%** expressed high confidence in their ability to detect one. H1 2026 supplied concrete illustrations of exactly that gap. On **4 May**, an attacker used an **encoded Morse-code message** to prompt-inject the **Grok AI model**, which directed an automated cryptocurrency bot to transfer roughly **$175,000** in tokens to the attacker's wallet. Around **80%** was later recovered, but the exposure itself, an AI system with enough standing permission to move real money on a single unverified instruction, is the point. On **12 May**, lawyers attempted the same technique against **Galileu**, a generative AI system used by a Brazilian labour court, hiding an instruction inside a routine filing. The attempt was detected and the lawyers sanctioned, but the method, not the outcome, is what should concern any organization feeding unfiltered external text into an AI system.

The **UK's National Cyber Security Centre** takes a more measured position than some of this vendor commentary, assessing that **fully automated, end-to-end advanced cyberattacks are unlikely before 2027**, with skilled human operators remaining in the loop even as individual elements of the attack chain continue to automate. We think that framing, not the more alarmist version, is the right one to build a defence posture around.

Guardrails

Two failure modes matter more here than any automation percentage.

False positives at scale are not harmless. An agent that suppresses genuine activity as noise, or escalates too aggressively, does not just waste analyst time; it creates a detection gap or trains a client to distrust the alerts they do receive. Autonomous high-impact action is the sharper risk. An agent that isolates a host, suspends an account, or blocks network traffic on a mistaken judgment can shut down legitimate business activity, a continuity incident, not a security success. High-stakes actions, host isolation, account termination, and production network changes must require human confirmation regardless of the agent's confidence score. That must be a designed control, not an assumption.

Where esentry Stands

We have adopted AI internally in two distinct ways, and we are deliberately not blending them into a single claim because they sit at different levels of maturity and risk.

The first is **Operational Automation for internal reporting**, which fast-tracks the production of weekly and monthly operational reports. This is workflow automation applied to reporting output: low-risk, already delivering time back to analysts and account teams for analysis and client engagement rather than report assembly.

The second is our **Agentic AI SOC capability**, launched this year and currently operating at partial autonomy. As the **COO's Desk** states elsewhere in this report, it was built with human sign-off required at every high-impact decision point, consistent with the guardrail principle above: we were not willing to caution the market about a risk we had not solved ourselves.

Narrative – What H1 Tells Us

The SOC operations picture from H1 is sustained pressure met with disciplined defence. The volume and variety of attack activity across web applications, endpoints, and network infrastructure reflect a threat environment that has industrialized, where attack tooling is cheap, automated, and running continuously against every accessible surface.

What held the line was integration between detection layers: endpoint telemetry feeding into network correlation, threat intelligence informing rule updates, and analyst expertise escalation to the cases automated correctly. The 11-minute median detection time is a product of that integration operating consistently under high volume.

The concern going into H2 is coverage, not capability. Environments with sparse logging, absent endpoint agents, or limited network visibility have detection gaps no analyst skill can compensate for. The incidents in the DFIR section of this report are, in several cases, the direct consequence of those gaps existing in production environments.

Threat Hunting

'The hunter in Darkness'

36
Hunt queries deployed

6
Hunt categories covered

1
Supply chain true positive

Hunt Coverage H1 2026

Threat hunting operations in H1 2026 were conducted under a hypothesis-led model targeting adversary behaviors that automated detection rules are specifically designed not to surface activity that blends with normal operations, uses trusted tooling, or leaves no signature-based artefact. **36 hunt queries** were deployed across **six technique** categories spanning the full attack lifecycle.

Most hunts confirmed the absence of targeted adversary behavior in monitored environments. Absence confirmed through active investigation is a different posture from absence assumed through silence. Every result was triaged and assessed against well-established baselines before any conclusion was drawn.

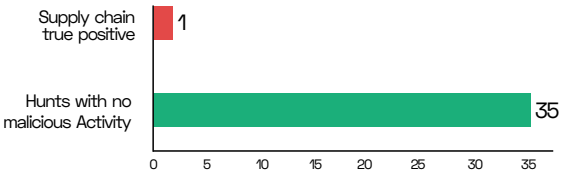

H1 2026 threat hunt query outcomes. 35 queries confirmed clean. 1 confirmed true positive

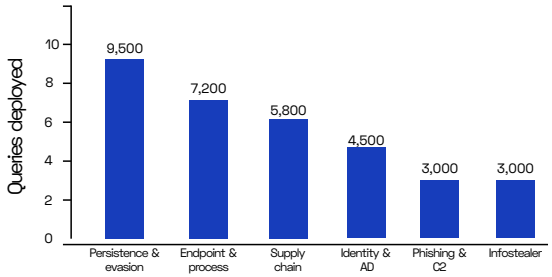

Hunt queries deployed by technique category, H1 2026.

Featured Hunt – When the Developer Workflow Became the Attack

The standout H1 threat hunting finding was not a vulnerability or phishing attack, but a trusted developer workflow being abused.An active Q2 2026 campaign embedded malicious instructions into legitimate repository files, allowing code to execute silently on developer machines and steal credentials, API tokens, and cloud secrets without triggering traditional security alerts. The activity appeared to be normal development work.The threat was identified through proactive hunting before any business impact occurred, enabling immediate credential rotation. The finding highlights a growing reality: the software development toolchain is now a critical attack surface that requires greater visibility and governance.

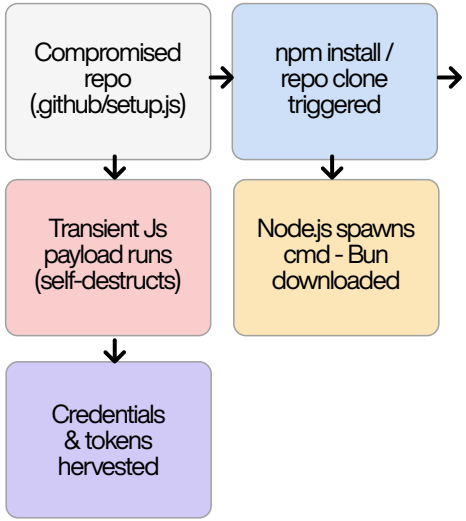

Miasma Worm execution chain. Detection requires behavioural hunting on process lineage — signature-based controls see nothing.

Hunt queries returned over **1,000 events** across monitored environments. Every result was triaged against **Miasma-specific behavioral indicators: process lineage**, **specific payload filenames**, and **DNS requests** to the Bun download URL. On one monitored endpoint, the full execution chain matched the associated setup.js file, which was confirmed malicious by **19 vendors** on Virus Total. Confirmed supply chain compromise identified through proactive hunting before any business impact was reported. No automated detection fired on this endpoint before the hunt ran.

Response: full repository integrity scan, rotation of all credentials and tokens accessible from the affected environment, CI/CD workflow history review, and environment-wide scope validation.

Browser Extensions: The Blind Spot on Every Endpoint

Over **10 suspicious** or unapproved browser extension detections were recorded in H1, with AI-assisted productivity extensions representing the fastest-growing category installed by employees without formal review or permission.

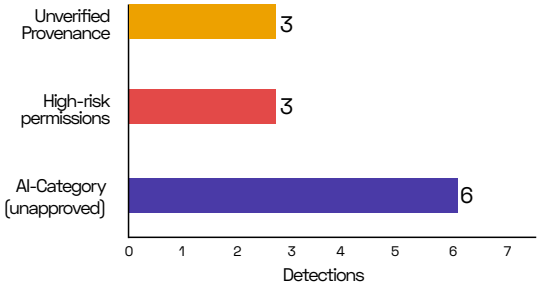

Fig B-4. Unapproved and suspicious browser extension detections by category, H1 2026.

A browser extension with web Request and tabs permissions can access everything the user does in their browser: form submissions, authentication tokens, session cookies, page content. Most enterprise endpoint tooling has no visibility at the browser session level. esentry maintains a threat-intelligence-driven extension blocklist and monitors suspicious extension-level behaviour in endpoint telemetry. The first line of defence, however, is an Acceptable Use Policy with enforcement, an approved extension list, a review process for new requests, and automated quarantine of unapproved installations on managed endpoints.

When the Inbox Became the Attack Surface – Trello and the OAuth Pivot

Business Email Compromise is no longer primarily an email problem. The most consequential shift in H1 2026 was the systematic move away from credential phishing toward identity compromise through OAuth consent abuse and session hijacking, which is a change that renders **SPF, DKIM**, and **DMARC** largely irrelevant as the primary defence.

A campaign active in Q1 and Q2 2026 used Trello's legitimate board invitation infrastructure as the delivery vehicle. The invitation email originates from **Atlassian's** own mail servers; SPF passes, **DKIM passes, DMARC passes**. The email lands in the primary inbox because it is technically a legitimate email from a legitimate platform. The threat is in what the embedded link delivers.

On click, the attack branches. One path route through an Adversary-in-the-Middle proxy that intercepts the post-MFA session token; the real user completes MFA, and the credential is captured in transit. The second path presents an OAuth consent dialog for persistent mailbox access scopes. No password required. One click grants a non-expiring refresh token that does not appear in email or login audit logs.

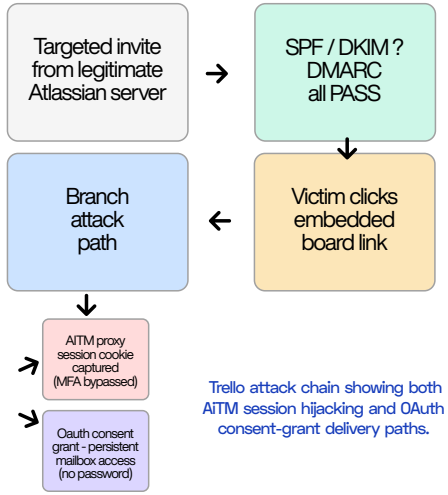

Trello attack chain showing both AiTM session hijacking and OAuth consent-grant delivery paths.

The same mechanism operates through **DocuSign notifications, SharePoint sharing alerts, Microsoft Teams invitations, and Google Workspace collaboration** prompts.

The platform changes; the technique does not. Detection requires behavioural signals: unexpected **SaaS** invitations from unrecognized organizations, OAuth consent grants requesting mail access scopes, mailbox rule creation shortly after an external invitation, and new OAuth applications appearing in the tenant consent list without a corresponding approval request.

The supply chain detection, browser extension findings, and the Trello campaign confirm the same structural reality: the attack surface has moved into the tools and platforms employees use as part of normal, trusted work. Developer workflows, browser sessions, and SaaS collaboration platforms are primary attack vectors in 2026. Organizations whose security posture is calibrated for perimeter defence and email-borne threats have a detection gap that is actively being exploited.

Threat Intelligence

"The Oracle of Horus"

Five Defining Trends of H1 2026

Ransomware settled at a new baseline.

Encryption alone stopped being the objective. Operators combined **double extortion**, **pre-encryption data theft**, **supply-chain delivery**, **cross-platform encryption**, and **deliberate destruction of forensic evidence**. We assess that global ransomware activity established a new operational baseline in early 2026 rather than returning to historical volumes; organizations should plan for sustained elevated attack volume, not a temporary spike. Emerging groups worth tracking include **Qilin**, **0APT**, **DragonForce**, **CipherForce**, **AiLock**, and **Handala**. **0APT** stood out for maintaining long-term access before encrypting, prioritizing data theft over immediate disruption.

H1 2026 Intelligence Snapshot

Intelligence Area	H1 2026 Observation
Primary Initial Access Vector	Identity compromise, verified, **MFA** -enabled, and **KYC**-bundled credential packages have overtaken traditional malware as the preferred route in.
Fastest Growing Threat	AI-enabled cybercrime, automated phishing, malware development, fake identity generation, and verification-bypass tooling.
Most Targeted Technologies	Software supply chains (**npm, PyPI, CI/CD pipelines**, developer tools) and unpatched internet-facing enterprise platforms.
Most Active Criminal Business Model	Ransomware-as-a-Service alongside data-broker extortion, ByteToBreach is a direct example of the latter.
Underground Economy	Verified and MFA-enabled identity packages now command premium pricing over raw stolen passwords; synthetic identities are the next iteration.
Primary Defensive Challenge	Detecting compromise that arrives through a trusted channel, a developer tool, a SaaS invitation, or a legitimate credential, rather than obvious malware.
Highest Risk Sectors (Nigeria)	Government (roughly one-third of documented incidents), Financial Services, Technology, Oil & Gas.
Strategic Assessment	Trust, not vulnerability exploitation, is now the primary surface of attack. Attackers are increasingly let in rather than breaking in.

THREAT INTELLIGENCE – "The Oracle of Horus'

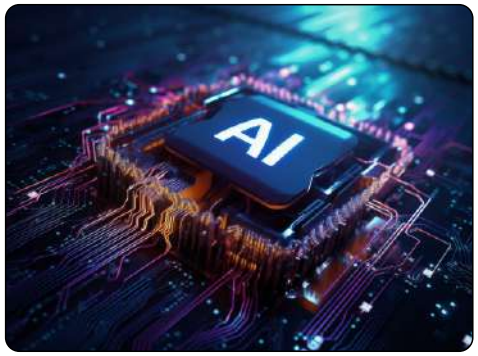

Software supply-chain attacks accelerated. Rather than attacking organizations directly, threat actors compromised the developer tools those organizations already trusted, **npm** packages, Python libraries, **CI/CD** pipelines, and open-source utilities. **TeamPCP's** campaign against **Trivy, LiteLLM, Bitwarden CLI, the Telnyx SDK, and Checkmarx KICS** is the clearest example: once developers updated an infected package, API keys, **Kubernetes** secrets, **SSH keys,** and cloud credentials were exfiltrated automatically. Full profile below.

AI became both a weapon and a target. Threat actors used AI to automate phishing, generate malware, write malicious code, fabricate identities, and conduct reconnaissance. At the same time, AI platforms themselves became targets worth stealing. **[Demand for premium AI accounts on criminal marketplaces reportedly rose sharply between Q4 2025 and Q1 2026, specific multiplier pending source confirmation.]** Emerging tradecraft in this space includes prompt poaching, memory poisoning, AI agent hijacking, and commercialization of unrestricted "dark" **LLMs**.

Credential theft replaced traditional malware as the preferred way in. Criminal marketplaces professionalized around selling verified accounts, MFA-enabled credentials, **KYC packages**, **session cookies, and browser fingerprints**, complete identity packages capable of bypassing modern authentication, not just passwords. **ByteToBreach's** own reliance on **infostealer logs**, detailed below, is a direct, documented example of this shift.

Synthetic identity fraud expanded the problem. Rather than stealing an existing identity, criminals increasingly assembled new ones, legitimate government identifiers, fabricated personal information, AI-generated documents, and forged verification artefacts. The absence of a real victim to report the fraud is what makes this category harder to detect than conventional identity theft.

Intelligence Analysis

Threat Actor Profile: TeamPCP

Attribute	Assessment
Assessment	HIGH RISK
Motivation	Financial
Primary Capabilities	Supply-chain compromise; credential theft; cloud exploitation; container compromise; CI/CD attacks;ransomware partnerships.
operational mode	Compromises trusted software ecosystems (developer tools and package registries) rather than attack victim organizations directly, maximizing downstream victim impact from a single point of compromise.

Threat Actor Profile: Payload Ransomware

Payload represents a technically mature ransomware family capable of targeting both Windows and VMware ESXi environments.

Attribute	Assessment
Platforms targeted	Windows, VMware ESXi
Cryptography	ChaCha20 encryption; Curve25519 key exchange
Extortion model	Double extortion
Anti-recovery capabilities	Anti-forensics; shadow copy deletion; ETW patching; backup destruction
Notable trait	Virtualization targeting significantly increases operational impact following compromise

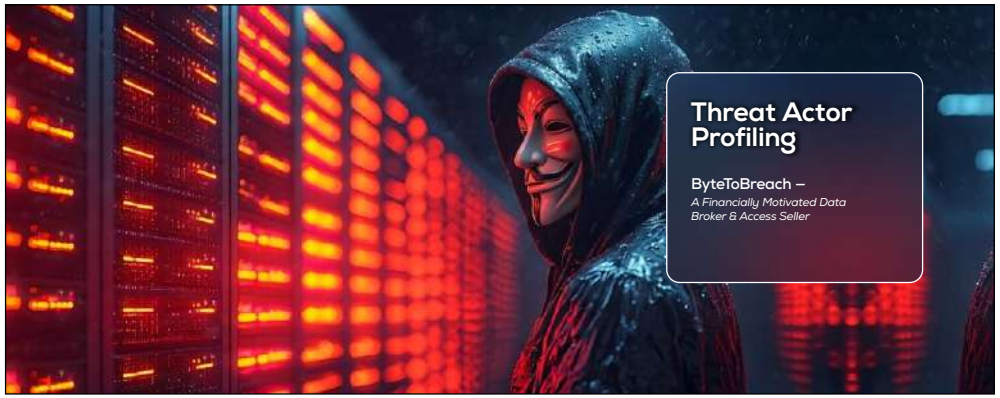

Threat Actor Profile: ByteToBreach – A Financially Motivated Data Broker & Access Seller

Threat Level	First Observed	Last Activity
HIGH	June 2025	March 2026

ByteToBreach is a financially motivated threat actor, likely an individual or small group, who has been actively selling stolen databases and network access on underground **cybercrime forums** since at least **June 2025**. The actor is prolific, claiming breaches across more than **35 organizations** in over **26 countries,** with financial services institutions accounting for the largest share of confirmed victims.

The **actor's modus operandi** centers on exploiting unpatched, internet-facing enterprise platforms (**GLPI, Oracle WebLogic, LMS** systems), extracting structured databases containing customer **PII, KYC records**, **credentials**, and **internal operational data**, and then monetizing these through dark web marketplaces or direct extortion. When organizations do not respond to ransom demands, **ByteToBreach** escalates to public shaming on their self-operated leak site, **bytetobreach.com**, and **forum postings**.

The **Nigerian DarkForums** thread advertising an alleged database from a major Nigerian bank represents a documented instance of the actor targeting Nigerian financial infrastructure. Given Nigeria's rapidly expanding fintech ecosystem and identified structural cybersecurity weaknesses, this profile is intended to equip **Nigerian banks**, **fintechs,** and **payment processors** with actionable intelligence.

Identity & Attribution

ByteToBreach operates under a single, consistent alias across **DarkForums (darkforums.su), Telegram, Instagram,** and its self-operated website. **OSINT** investigation traced the alias to a **Greek-speaking** individual based in **Thessaloniki**, **Greece,** using a username pattern containing the **Greek letter 'η'** across multiple platforms. The actor presents themselves publicly as a penetration tester and security consultant via **bytetobreach.com** – a tactic common among extortion actors that allows them to approach victims as though offering paid remediation.

Attribution confidence is MODERATE-HIGH (a single operator or very small group). The actor demonstrates consistent operational security habits, including use of **Tuta encrypted email**, Signal, **Tails OS**, and **Kali Linux,** while simultaneously leaving exploitable digital footprints across public social media, freelancer sites, and data breach databases.

THREAT INTELLIGENCE – "The Oracle of Horus'

Motivation & Model

ByteToBreach is purely financially motivated. Unlike hacktivist groups or nation-state actors, the actor selects targets based on data value and access potential rather than political affiliation or geography. The operation follows a structured monetization model:

1. Initial breach and data exfiltration from internet-exposed enterprise systems.
2. Direct, private extortion attempt — the actor contacts the victim organization with a demand.
3. Escalation to a public dark web forum listing if no response is received within a set window.
4. Concurrent or follow-up publication on bytetobreach.com for victim-shaming and broader exposure.
5. Ongoing sale of access ('reverse shell') to the highest bidder, sometimes independently of the extortion outcome.

Known Aliases & Contact Infrastructure

- **Primary alias: ByteToBreach**
- **Dark web forum: DarkForums (darkforums.su)**
- **Website / data-leak site: bytetobreach.com (WordPress)**
- **Email: Bytetobreach@tuta.com**
- **Encrypted messaging: Signal handle 'Bytetobreach'**
- **Alternate Telegram handles: CvHNWwEG, iɳesslopez**
- **Linked real-world identity (OSINT): individual named Anastasis from Thessaloniki, Greece**

Tactics, Techniques & Procedures (MITRE ATT&CK)

Mapped from confirmed incidents across the actor's documented victim set, spanning telecom/managed infrastructure, regional banking, aviation, and enterprise technology sector targets.

Tactic	Technique	Observed Behaviour
Resource Development	T1583.003 — Virtual Private Server	Rented ~20 VPS nodes across EU countries to parallelise SQL injection queries and anonymise exfiltration traffic
Resource Development	T1583.006 — Web Services (WordPress)	Operates bytetobreach.com as a victim-shaming and data-leak site disguised as a pen-testing firm
Initial Access	T1190 — Exploit Public-Facing Application	SQL injection against internet-exposed GLPI, Oracle WebLogic,LMS/ticketing platforms; time-based blind injection observed in a confirmed victim case
Initial Access	T1078 — Valid Accounts / Stolen Credentials	Uses infostealer logs and phishing-derived credentials to authenticate to exposed services
Execution	T1059 — Command and Scripting Interpreter	Advertises reverse-shell access post-exploitation; shell activity implied by forum posts advertising enterprise technology/telecom access
Discovery	T1082 — System Information Discovery	Maps internal ITSM platforms, configuration items, and network topology via compromised service-desk access
Collection	T1213 — Data from Information Repositories	Extracts structured databases from ITSM, GLPI, LMS, banking core systems, and document stores
Exfiltration	T1048 — Exfiltration Over Alternative Protocol	Uses VPS relay infrastructure; slow time-based extraction (~10 days in one confirmed case) to evade rate-limit detection
Impact	T1657 — Financial Theft / Extortion	Contacts victim directly for ransom; escalates to public forum posting and leak-site exposure if ignored
Impact	T1486 — Data Encrypted for Impact (attempted)	In one confirmed banking-sector case, attempted leverage via hash-cracking and threatened decryption disclosure

Deep Dive: SQL Injection at Scale

The most technically significant documented **TTP** is the actor's use of time-based blind SQL injection against a confirmed victim organization. After identifying an unpatched **GLPI** instance exposed directly to the internet, the actor executed slow, time-based queries over a period of approximately ten days. To parallelize the extraction and remain below rate-limiting thresholds, approximately **20 VPS servers** rented in multiple European countries were used, rotating source IPs and distributing query load. This approach extracted roughly **10,000 password hashes, configuration data, credentials for downstream customer organizations**, and **service management tickets**, a sophisticated, operationally patient approach that contrasts with typical opportunistic smash-and-grab intrusions.

Deep Dive: Oracle WebLogic Exploitation

In another confirmed case, the actor claimed to have leveraged a known Oracle WebLogic vulnerability to gain an initial foothold at a regional bank.

Once inside, the focus shifted to banking data and government records held within the institution, totaling **2.2 GB of exfiltrated data**. The actor then attempted to use decryption capabilities as leverage, offering to sell back access or withhold publication.

Sector Targeting Summary

Sector	Risk Level
Financial Services (Banks, Fintech, Insurance)	HIGHEST
Telecoms & Infrastructure	HIGH
IT Services & Data Centers	HIGHEST
Airlines & Logistics	MEDIUM
Government / Public Sector	MEDIUM
Education & Healthcare	MEDIUM
Real Estate & Diversified	LOW-MED

Indicators of Compromise (IOCs)

Type	Indicator	Notes
Domain	bytetobreach.com	Actor-operated victim-shaming / leak site
Email	Bytetobreach@tuta.com	Contact email used in forum posts
Signal Handle	Bytetobreach	Encrypted messaging contact
Dark Forum	darkforums.su (DarkForums)	Primary marketplace for access and data sales
Alias	ByteToBreach	Consistent handle across all underground platforms
Alt Handle	iηesslopez (Greek η character)	Secondary Telegram username; Greek-language clue
Infrastructure	~20 VPS nodes, European providers	Used to parallelize SQL extraction in a confirmed case
CVE Exploited	Oracle WebLogic (unspecified)	Confirmed banking-sector compromise
CVE Category	SQL Injection in GLPI (unpatched)	Time-based blind injection technique, confirmed case

Dark Web Presence & Forum Activity

ByteToBreach maintains an active presence on DarkForums, one of the most active English-language cybercrime forums accessible via the standard web. Forum posts follow a template: organization identity, sector, country, data size, data types, price, and contact method. The actor also advertises reverse-shell access separately from database dumps, indicating they maintain persistent access to some compromised environments even after selling the initial data.

The self-operated bytetobreach.com website lists victims with company logos and fabricated 'client testimonials', framing the hacker as someone victims eventually hired to protect them. This tactic serves dual purposes: maintaining a veneer of legitimacy and applying additional psychological pressure on victim organizations to negotiate.

TeamPCP: GitHub Internal Repository Breach

On 20 May 2026, GitHub disclosed that attackers had gained unauthorized access to internal repositories after compromising an employee device through a poisoned VS Code extension. The attack was reportedly detected and contained quickly, with GitHub removing the malicious extension, isolating the affected endpoint, and initiating incident response procedures.

TeamPCP later claimed responsibility for the intrusion and allegedly offered approximately 3,800 stolen repositories for sale on underground forums for more than **$50,000.** The group reportedly stated that the operation was not ransomware-driven but instead focused on selling or leaking the stolen data. GitHub acknowledged that the attacker's claims regarding the number of repositories were 'directionally consistent' with the company's ongoing investigation.

This incident highlights the growing risk posed by software supply-chain attacks targeting developer environments rather than core production infrastructure — a pattern this assessment has separately tracked in TeamPCP's compromise of **Trivy, LiteLLM, Bitwarden CLI, the Telnyx SDK, Checkmarx KICS,** and numerous **npm** packages.

How the Attack Happened

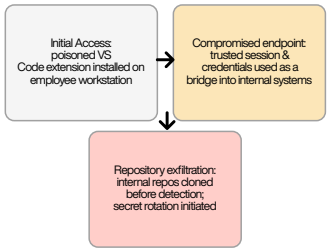

GitHub internal repository breach: attack chain.

How the Attack Happened

The attack originated from a malicious VS Code extension installed on a GitHub employee's machine. Once activated, the extension reportedly executed hidden malicious code capable of interacting with the developer's environment and existing authentication context.

VS Code extensions are particularly dangerous because they often receive broad permission, including:

- File system access
- Network connectivity
- Terminal execution
- Workspace inspection
- Credential interaction

- Workspace inspection
- Credential interaction

This effectively allows malicious extensions to behave with elevated privileges inside developer environments.

Compromised Developer Endpoint

After infecting the employee workstation, the attackers leveraged the trusted session and existing credentials to access GitHub's internal repositories. Rather than directly breaching GitHub's infrastructure externally, the attackers used the developer machine as a trusted bridge into internal systems. This technique reflects a modern software supply-chain compromise model where attackers bypass hardened perimeter defences by targeting:

- Trusted developer tools
- End-user devices
- Extension marketplaces
- Authentication tokens
- CI/CD secrets

Repository Exfiltration

The attackers then cloned or exfiltrated internal repositories before detection occurred. GitHub stated that it immediately initiated secret rotation and forensic review after identifying the compromise. At the time of reporting, GitHub stated there was no evidence that attackers modified production systems or inserted malicious code into public repositories.

Criminal Collaboration: TeamPCP and LAPSUS$

An additional development that significantly elevates the severity of this incident is the reported collaboration between TeamPCP and the notorious cybercriminal group LAPSUS$. According to a widely circulated dark web monitoring alert, **TeamPCP** and **LAPSUS$** jointly advertised the sale of GitHub internal repositories allegedly stolen during the breach.

The forum post claimed:

- Access to approximately **4,000** private repositories
- Internal GitHub source code and organizational data
- A sale price ranging from **$50,000** to **$95,000**
- Intent to leak the data publicly if no buyer emerged
- A 'no ransom' motivation, emphasizing monetization and exposure rather than extortion

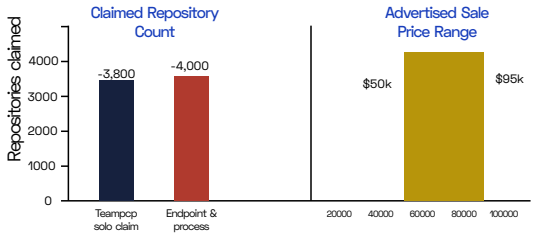

Claimed repository count and advertised monetization range

4.1 Why the LAPSUS$ Association Matters

Increased Credibility of the Threat

LAPSUS$ has historically targeted major technology companies through social engineering, insider compromise, credential theft, and developer-platform intrusion, with previously confirmed impact against organizations including **Microsoft, NVIDIA, Uber, Samsung,** and **Okta**. Their involvement — whether operational, collaborative, or purely reputational — adds credibility to the breach claims and increases the likelihood that the stolen data is authentic and potentially sensitive.

LAPSUS$ partnering with TeamPCP for data monetization suggests:

- Established cybercriminal market - place relationships
- Coordinated leak operations
- Shared infrastructure or access-broker relationships
- Expanded visibility across underground forums

This amplifies both reputational and operational risk for GitHub and any potentially affected downstream systems, and by extension for any organization that depends on GitHub-hosted infrastructure, actions, or extensions as part of its own software supply chain.

Major Cyber Incidents and Threats Observed in Nigeria, H1 2026

Nigeria did not experience a single nationally disruptive event on the scale of the global supply chain compromises described above. Still, intelligence indicates persistent, cross-sector targeting, sharpened by a confirmed, named data broker actively advertising Nigerian financial-sector data.

Sector Observations

Sector	H1 2026 Pattern
Government	The most targeted sector nationally, roughly one-third of documented incidents, reflecting sustained interest in public-sector infrastructure and operational disruption
Financial Services	Phishing, credential theft, BEC, and ransomware activity, sharpened by confirmed ByteToBreach targeting
Technology	Increasing exposure through software supply-chain attacks and cloud infrastructure compromise
Oil & Gas	Remains attractive due to operational importance and economic impact

Statistical Snapshot: Nigeria's Breach and Attack Exposure

Quarterly breach-tracking data places Nigeria among the more heavily exposed markets globally in early 2026: **281,500 accounts** tied to Nigerian users were newly breached in Q1 alone, ranking the country 34th most breached worldwide for the quarter. Cumulatively, Nigeria has recorded **24.1 million** compromised accounts since **2004**, the third highest in **Sub-Saharan Africa**.

Nigeria - Cumulative Data Exposure Since 2004

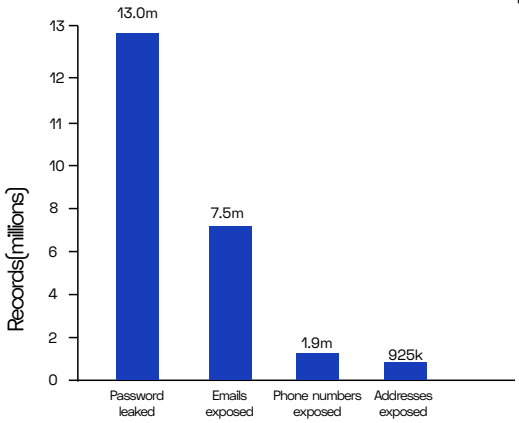

Figure 1 — Nigeria: cumulative exposure of large-scale record categories since 2004.

A meaningful share of this exposure involves highly sensitive, hard-to-rotate data rather than passwords alone. Roughly **3,900 Social Security-related records** and **1,600** payment card details tied to Nigerian users have been identified in circulation, alongside the volumes shown above.

Nigeria - Highly Sensitive Records Exposed

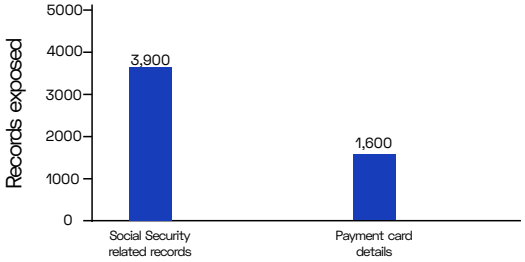

Fig 2 — Nigeria: highly sensitive record categories exposed.

An estimated **54% of breached Nigerian** users face heightened risk of account takeover, identity theft, or extortion, consistent with tracking showing roughly **1 in 10 Nigerians** affected by a breach at some point. This is compounding against accelerating global volume: breached accounts rose **22%** quarter-on-quarter worldwide in Q1 2026 and are estimated to have roughly tripled year-on-year.

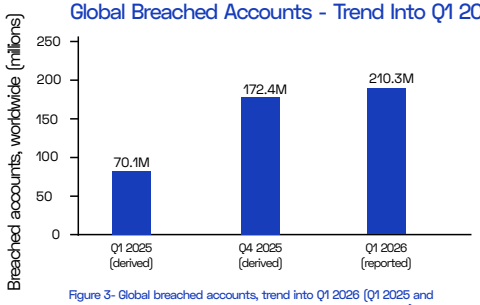

Figure 3- Global breached accounts, trend into Q1 2026 (Q1 2025 and Q4 2025 figures derived from reported YoY/QoQ growth rates).

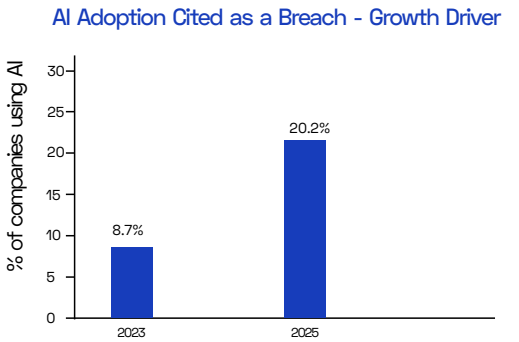

Figure 4 — Enterprise AI adoption, cited as a contributing driver of breach-volume growth.

Industry commentary attributes part of this acceleration to the pace of enterprise AI adoption: the share of companies using AI in their operations is estimated to have grown from **8.7% in 2023 to 20.2% in 2025**, with AI-driven systems logging larger volumes of user data across a growing number of connected systems — widening the exposure surface even as it improves operational efficiency.

At the network level, Nigeria's national cyberattack volume is estimated at over **4,000–4,700** attacks per week — a figure corroborated at the individual-organization level, with the country's own data protection regulator reporting over **1,500 attack attempts** against its own network in a single week.

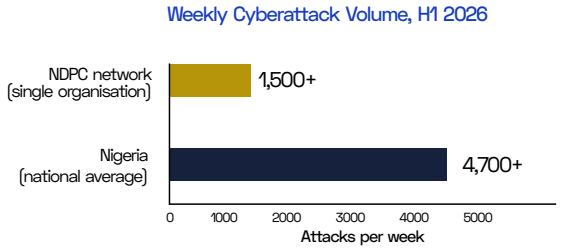

Figure 5 — Weekly cyberattack volume: national average vs. a single regulator's network.

Regulatory Alert: Coordinated Threat Activity Against Financial and Digital Infrastructure

Nigeria's data protection regulator issued a formal advisory in April 2026 after a technical assessment uncovered coordinated activity by **"shadowy threat actors"** against key national systems, flagging banking, payment platforms, telecommunications, cloud infrastructure, and public-sector digital services as increasingly vulnerable.

Directed Action	
Governance	Appoint trained, certified Data Protection Officers; implement privacy policies and information security standards; conduct Data Privacy Impact Assessments on critical systems
Access & Monitoring	Deploy identity and access controls, including mandatory MFA; enable real-time monitoring, logging, and threat detection; encrypt and secure credential management
Resilience	Conduct regular vulnerability assessments and penetration testing; test backup, recovery, and resilience procedures

The advisory landed amid an active regulatory investigation into an alleged breach involving a major payment-services provider and a Nigerian bank, and alongside a broader enforcement push: the regulator is separately probing more than **1,300 companies,** predominantly financial-sector and government entities, over data privacy violations, and has previously fined a major pay-TV operator over NGN 766 million for unlawful data processing. That enforcement posture materially raises the compliance stakes for any Nigerian financial institution found with inadequate safeguards at the time of an incident.

Escalating DDoS Threat to Critical Infrastructure

Nigeria's national CERT separately warned of a sustained, intensifying wave of DDoS attacks against government and private-sector infrastructure, increasingly complex campaigns blending volumetric floods, protocol-based exploits, application-layer attacks mimicking legitimate traffic, and DNS/NTP/Memcached reflection and amplification. Attackers were observed exploiting long-known, unpatched vulnerabilities, including **Log4Shell and flaws in Citrix and Drupal,** to recruit servers, endpoints, and IoT devices into botnets.

The advisory's sharpest point: **DDoS** activity is in some cases being used as a diversionary tactic to mask more severe intrusions, ransomware deployment, or data exfiltration. A DDoS event against any client environment should trigger a check for concurrent, quieter compromise, not just an availability-focused response.

Confirmed Targeting: ByteToBreach and the Wider Nigerian Breach Pattern

ByteToBreach has advertised what it claims is a database from a major Nigerian bank on DarkForums, unconfirmed by the institution at the time of publication, but consistent with the actor's established pattern (see full profile above). The Nigerian financial institutions are a natural target for this actor: the sector processed over **108 billion** mobile transactions worth **$1.68 trillion** in 2024, digital fraud rose **468%** between **2023** and the first nine months of **2024,** and many banks and fintechs rely on the exact internet-exposed, third-party-integrated infrastructure ByteToBreach targets.

This listing isn't isolated. A national companies-registry agency was forced to temporarily suspend its registration portal after ByteToBreach claimed exfiltration of roughly **25 million files (~750 GB),** demonstrating willingness to escalate from financial targeting to core government registry infrastructure. Other organizations named in alleged or confirmed activity during the same window include a major payment-services provider, a national law-enforcement agency, a public university, and a consumer lending platform, indicating the threat is actively probing across financial, regulatory, law-enforcement, and public-education infrastructure concurrently, not confined to one sector.

Timeline of Confirmed and Alleged Incidents, H1 2026

Timing	Sector / System	Status	Notes
Apr 2026	Payment services provider & a major bank	Under regulatory investigation	Triggered the regulator's formal advisory
Apr 2026	National Companies Registry Agency	Claimed by ByteToBreach	~25M files / ~750GB allegedly exfiltrated; portal temporarily suspended
Apr 2026	National law-enforcement agency	Alleged	Named in the same wave of public breach claims
Apr 2026	Public university	Alleged	Named in the same wave
Apr 2026	Consumer lending platform	Alleged, unconfirmed	Sample data posted as proof of control
Ongoing	Major Nigerian bank (via DarkForums)	Alleged, unconfirmed	Primary subject of the ByteToBreach profile above

Two things stand out here. Several of these are regulators, registries, or law-enforcement bodies rather than commercial targets, suggesting attackers value the downstream leverage of government-held data as much as direct extortion. And the clustering around April 2026, coinciding with the regulator's own advisory, suggests either a coordinated campaign window or that regulatory scrutiny surfaced incidents that had been running longer than the public timeline shows.

Active Threat Actors Against Nigerian Organizations

Threat Actor	Type / Notes
XForum Bot	Automated/bot-driven; highest-frequency named actor against Nigerian targets
AlQe3Qa3	Consistent targeting pattern across multiple sectors
DitzzXploit	Exploitation-focused methodology
XYZEAZ	Active across government and business sectors
Incransom	Ransomware/extortion, encryption-based
Killsec	Ransomware/extortion, encryption-based
ByteToBreach	Data broker/access seller, full profile above

Security Testing

"The Sword of Horus "

SECURITY TESTING: "The Sword of Horus "

Nigerian financial infrastructure, payment systems, and **digital businesses** are under active exploitation right now. The attackers are not sophisticated. They are methodical. They probe for the one credential that was reused, the one API endpoint left without authorization checks, the one administrative account that accesses more than it should. When they find it, they move. They extract. They disappear.

We spent H1 2026 testing exactly what those attackers are finding. Across **175** engagements, the same vulnerabilities kept appearing. The same patterns. The same exploitability. But we also found something else: organizations that got security controls demonstrated measurably greater resilience. The difference was not budget or sophistication. It was architectural discipline and operational rigor.

This section documents what we proved can be broken, how organizations are failing, and which control decisions genuinely matter against attackers already hunting for these weaknesses.

Engagement Portfolio H1 2026

Engagement Type	Count	Avg. Success Rate
Targeted Penetration Tests	170	100%
Social Engineering Campaigns	4	100%
Cloud Security Assessments	1	100%

Top Vulnerabilities Observed – H1 2026

Vulnerability	Description	Why It Keeps Appearing	Observed Frequency
Broken Function Level Authorization (BFLA)	Users access functions beyond intended privileges.	Inconsistent backend authorization.	35–45%
Broken Object Level Authorization (BOLA/ IDOR)	Ownership validation is missing.	Trust in client-supplied identifiers.	45–55%
Hardcoded Secrets	Secrets embedded in code or apps.	Weak secrets management.	30–40%
Client-Side Cryptographic Failures	Keys exposed in client-side code.	Poor cryptographic architecture.	40–50%
Excessive Data Exposure	APIs return unnecessarily sensitive data.	Lack of backend data minimization.	60–70%

Why These Patterns Persist -Architectural Lessons from Recurring Vulnerabilities

The vulnerabilities identified reflect architectural decisions made early in the software development lifecycle. Broken access control, excessive data exposure, hardcoded secrets, and cryptographic failures appear repeatedly because they stem from design-stage assumptions rather than implementation oversights.

Modern applications increasingly rely on APIs, microservices, and client-side frameworks. When secure design principles are not consistently applied across this distributed architecture, the attack surface expands dramatically. Applications built without authorization logic at every API endpoint, without secret-management infrastructure, or without proper data classification and minimization principles will expose these vulnerabilities regardless of code review rigor or input validation strength.

Secure-by-design approaches require treating authorization as a first-class requirement, not an afterthought. Backend authorization enforcement must be assumed necessary at every function, not just at the perimeter. Secrets management must be centralized and mandatory. Data returned from APIs must be filtered to business-necessary fields only. These are not deployment problems to solve with WAF rules or API gateways. They are architectural problems that require discipline during the earliest design phases.

Internal vs External Engagement Findings

External engagements continued to expose internet-facing weaknesses such as outdated software, exposed services, insecure API endpoints, authentication weaknesses, and security misconfigurations. These issues generally represented the organization's initial attack surface and often required relatively little effort to identify. Misconfigured web applications, excessive information disclosure, and externally accessible management interfaces remained recurring themes across engagements.

Internal engagements presented a different picture once an initial foothold had been established. Weak network segmentation, legacy protocols, excessive privileges, stale accounts, and Active Directory weaknesses frequently enabled privilege escalation and lateral movement. These findings reinforce that while perimeter security continues to improve, many organizations still require stronger internal security controls to limit attacker progression.

Common Organizational Errors

The most common weaknesses observed were operational rather than purely technical. Organizations frequently struggled with environment segregation, privileged access governance, asset inventory management, patch management, and continuous validation of deployed security controls. These gaps often created unnecessary attack paths despite significant investments in security technologies.

Governance and ownership also remained recurring challenges. Shadow IT, forgotten systems, delayed remediation, inconsistent change management, and insufficient security testing before production deployments were regularly encountered. Organizations with mature governance processes consistently demonstrated fewer high-risk findings than those relying primarily on technical controls.

Notable Engagement Highlight: Authorization Failure and Business Impact

Situation and Initial Indication:

During one engagement in H1 2026, offensive security testing of a financial services application identified a critical authorization weakness affecting sensitive business functions.

Technical Findings

Testing identified that the application enforced authentication correctly but failed to implement consistent authorization validation across all API endpoints. An authenticated user with valid credentials could access functions and data far beyond their assigned role through direct manipulation of request parameters. The weakness allowed unauthorized access to sensitive operational information without requiring additional authentication or triggering any

Exploitation Complexity:

Although exploitation required minimal technical complexity (parameter modification and direct API calls), the potential business impact was significant. Sensitive operational information accessible through this vector included transaction data, customer information, and system configuration details that could support follow-on attacks.

Remediation and Validation

The organization implemented comprehensive authorization checks enforcing role-based access control across every application function. Following remediation, additional validation testing confirmed that the strengthened access control architecture significantly reduced the organization's overall exposure to this attack class.

Lessons for Other Organizations:

This engagement underscores a critical principle: authentication and authorization are distinct security functions. An organization may have strong authentication controls (MFA, password policies, login monitoring) yet remain vulnerable if authorization logic is absent or inconsistent. The most effective remediation requires that authorization be enforced at the backend layer of every API endpoint and business function, with the assumption that client-side controls cannot be trusted.

The first half of 2026 reinforced that fundamental security weaknesses continue to be a greater risk than sophisticated attack techniques. Access control failures, API security weaknesses, exposed secrets, and governance deficiencies remained the most consistent findings across engagements. These trends suggest that secure architecture and operational discipline continue to require greater attention across many organizations.

At the same time, positive progress was observed amongst organizations that have embraced secure development practices, stronger identity security, and continuous monitoring. As organizations move into the second half of 2026, continued investment in secure-by-design principles, proactive governance, and continuous validation of security controls will be critical to reducing recurring security risks.

What Organizations Must Prioritize in H2 2026

Do not wait for the next penetration test or incident to act on authorization. Across 175 H1 engagements, broken authorization appeared in nearly every assessment. Threat actors are exploiting this weakness in active breaches across Nigerian and African financial infrastructure right now.

Your immediate priorities should be to enforce authorization logic at the backend on every API endpoint. Centralize secrets management before any code reaches production. Implement data minimization as an architectural requirement, not an option. Validate these controls continuously through testing and monitoring.

Authentication alone will not protect you in H2. Authorization enforcement will.


DFIR: Digital Forensics &
Incident Response
The Scales of Maat

The Half in Numbers

3
Major engagements

0
Automated detections fired

3
Sectors covered

NGN 1B+
Aggregate fraud value

Three major **DFIR** engagements in H1 2026 across the financial technology, payment infrastructure, and core banking sectors. Collective confirmed fraudulent activity exceeded NGN 1 billion. Every incident was discovered through a manual process, routine account reconciliation, a third-party support ticket, or a national settlement system notification after fraud had been concluded.

The number that defines **H1 DFIR** is not the fraud value managed. It is the zero next to automated detections fired. Three separate engagements, three different attack mechanisms, three different sectors, and not one control in any affected environment detected anything while the attack was in progress.

The Structural Pattern

The technical specifics of each engagement were distinct. The structural failures were not. Three observations are held consistently across all three cases.

In every engagement, the attacker operated through something the environment already trusted. Valid API credentials used with a whitelisted IP. Authenticated sessions pivoted direct API calls from an unexpected origin. A compromised administrator account with legitimate access to the systems it manipulated. At the layers where controls were checking, the activity looked authorized. The anomaly was only visible at layers where no controls existed.

In every engagement, the logs that would have resolved the most critical investigative questions were absent. Network-level logging is not enabled in one environment where a fraud script executed inside the client's own cloud infrastructure and the responsible workload cannot be identified retrospectively. Authentication records without source IP retention in another session continuity analysis were impossible. In a third, Windows Event Logs were deliberately cleared by the attacker before forensic collection began, permanently destroying evidence from the compromise period.

In every engagement, a third party was either the entry point or a significant investigative blind spot. A payment provider processed hundreds of fraudulent transactions over several hours with no velocity alert and no proactive notification to the affected client. An abused internal access control workflow introduced a remote access tool that became the initial foothold. Credentials managed by infrastructure automation were accessible to any workload in the same environment namespace.

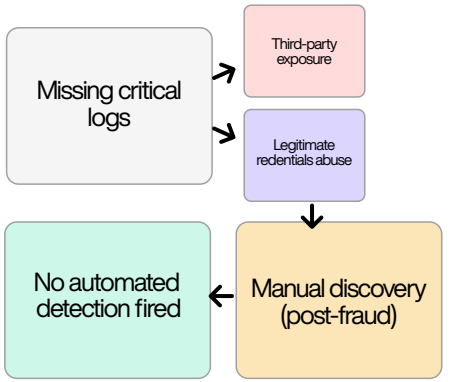

Structural pattern consistent across all three H1 2026 DFIR engagements

What This Means for Security Leadership

The incidents in this section are representative of the threat environment currently operating against Nigerian and African financial infrastructure. The common threat is not attack sophistication; it is the gap between what the environment could see and what was happening.

The most operationally important question these cases raise for CISOs and CIOs is not whether perimeter controls are sufficient. It is whether your organization would know, in real time, if its own credentials were being used against it. Whether a script running inside your own infrastructure would produce any log entry that anyone would ever see, or your payment provider would notify you proactively if something anomalous was occurring on your account.

The gap between "we have security controls" and "we would detect this" is wider than most organizations appreciate. These findings are a precise map of where that gap sits.

Security Engineering

Builders of the Eye

Security Engineering: Builders of the Eye

Every finding, alert, and vulnerability identified in the paths of this report eventually raises the same question: **what gets built to close it?**

This section covers the work for H1 **2026;** the security implementations completed in response to identified gaps across esentry's customer base, the threats and operational weaknesses each was designed to address, the architectural shortcomings observed during delivery, the lessons drawn from the deployment process itself, and a forward-looking set of recommendations for H2.

Security engineering effort over the period spanned three layers of customer defence: security monitoring and incident detection, endpoint detection and response, and web application protection. Nine customer engagements were completed across these three areas, each initiated to close a specific, identified gap rather than as a routine technology refresh.

9

Gaps Closed

3

Layers Of Defence Strengthened

Each engagement was undertaken in response to a specific visibility or control gap identified in a customer environment. The table below summarizes the solutions deployed, the gap each was intended to close, and the number of customers served.

Focus Area	Solutions Deployed	Gap / Threat Addressed	Customers Served
Security Monitoring & Incident Detection	LevelBlue USM Anywhere, Rapid7 Incident Command	Visibility gaps and log fragmentation; limited attack surface management and threat intelligence capability	4
Endpoint Detection & Response	CrowdStrike, SentinelOne	Absence of endpoint detection and response capability; treated as foundational hardening	3
Web Application Protection	Apptrana, FortiWEB	Unprotected customer-facing and business-critical web applications; limited attack visibility	2

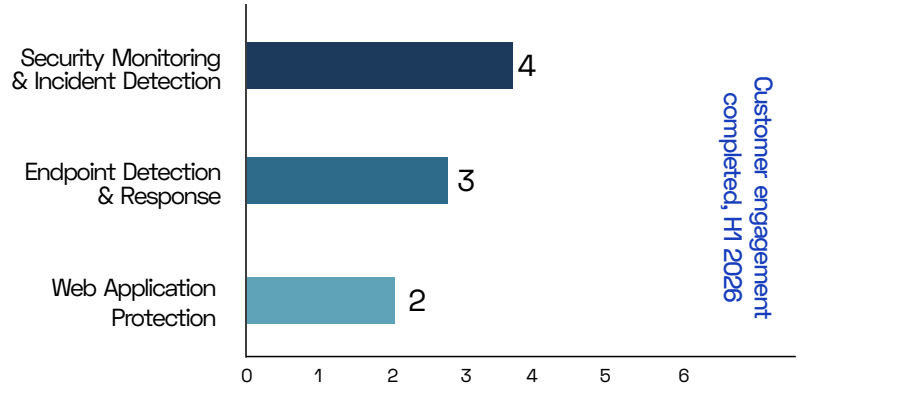

Security Engineering: Builders of the Eye

Security monitoring and incident detection accounted for the largest share of H1 deployments. This reflects two distinct drivers rather than one trend: some engagements were initiated in direct response to attacks already observed in a customer environment, while others formed part of a longer-term effort to establish and maintain baseline visibility. Endpoint detection and response deployments were approached consistently as foundational hardening rather than a reactive measure, consistent with the broader shift toward endpoint protection as a minimum-security requirement discussed elsewhere in this report.

Common Architectural Weaknesses

Two weaknesses recurred across customer environments during delivery. Neither was specific to a single deployment, and both pointed to gaps in ongoing security operations rather than one-off configuration errors.

Tooling decays after deployment.

EDR agents, SIEM agents, and other integrated event sources are frequently left unmaintained after initial deployment, creating visibility gaps that erode the value of the tooling over time.

Incomplete SIEM source coverage.

Customers consistently integrate fewer event sources into their SIEM than their environment requires, narrowing monitoring coverage and limiting detection effectiveness from the outset.

Taken together, these patterns indicate that more persistent risk is not the absence of security tooling but the absence of sustained ownership of it once deployed.

Deployment Experience and Lessons Learned

Most H1 deployments were completed rapidly, supported by close collaboration with customer technical teams, allowing value to be demonstrated within a short timeframe. Where delays occurred, the cause was consistently on the customer side: internal approval processes and, in some cases, infrastructure availability, rather than any factor within Security Engineering's control.

For H2, this points to a change in approach rather than a change in tooling: setting out deployment requirements and realistic timelines with customers upfront, so that internal approvals and infrastructure provisioning are secured before they become a bottleneck to delivery.

Recommendations for H2

The following recommendations are drawn directly from patterns observed during H1 engineering delivery. Recommendations without a direct evidentiary basis in this period's data have been excluded.

Maintain deployed tooling; do not just deploy it.

Organizations should establish scheduled health checks for already-deployed EDR and SIEM agents and integrations, given that unmaintained tooling was the most consistently observed weakness across H1 engagements.

Expand SIEM source coverage.

Organizations should periodically audit and expand the event sources feeding their SIEM, as incomplete source integration recurred across customer environments and directly limits detection coverage.

Treat endpoint protection as baseline control.

Organizations without any form of endpoint detection and response should prioritize deployment, consistent with H1 engagements that treated EDR as foundational rather than optional hardening.

Global Threat Perspectives: Darktrace

Information presented here is from Darktrace Annual Threat Report and recent Darktrace findings in collaboration with esentry

The Eye of Horus sees most clearly when it looks beyond its own boundaries. In this section, we incorporate perspectives from a respected voice in global cybersecurity.

What Darktrace Observed in Africa

StealC infostealer activity.

Customers across the **EMEA** and **AMS** regions, including several in Africa, were affected by the StealC infostealer in late 2025 and early 2026.

Hola VPN abuse.

A campaign identified in February and March 2026 saw the use of **Hola VPN** followed by malicious payload deployment and, in some cases, escalation to crypto mining. Africa was particularly targeted, forming 20% of all identified victims across the wider customer base.

Phishing-to-cryptomining chains.

Phishing attacks leading to malware downloads and subsequent cryptocurrency mining were also observed affecting African customers in early 2026.

Unusual .biz domain connectivity.

Their threat research team is tracking a pattern of unusual connectivity, outgoing HTTP POST requests to external .biz domains, seen since April 2026 and significantly affecting African customers. This remains active tracking rather than a confirmed finding at the time of publication.

None of this happened in isolation. Each of these H1 2026 patterns sits inside a longer regional trend, and Darktrace's broader Africa dataset, covering 2024 to 2025, helps explain why. The infostealer activity, VPN abuse, and payload delivery chains observed this period are variations on a threat landscape that has looked broadly similar in Africa for several years.

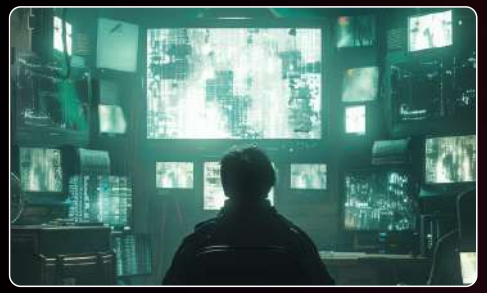

The Regional Backdrop: Africa, 2024 to 2025

As one of the world's largest cybersecurity vendors, Darktrace's visibility spans dozens of countries and industries. Turning that global lens onto Africa, their regional analysis for 2024 to 2025, the period immediately preceding this report, shows a threat landscape dominated by network-based activity and ransomware, set out below.

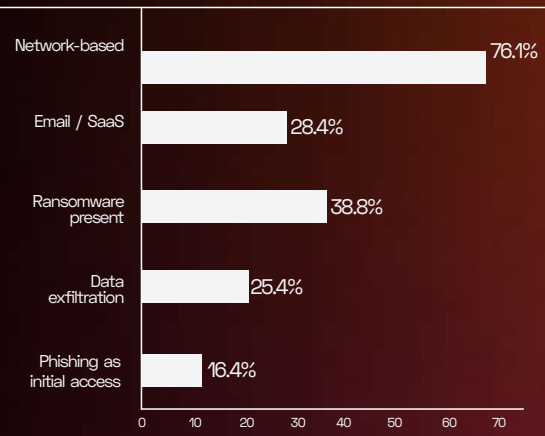

Share of Darktrace-detected incidents, Africa region, 2024-2025

The data also points to a consistent behavioural pattern across these incidents: attackers blending legitimate administrative tools, including **RDP, SMB, WinRM, PowerShell,** PsExec and VPN infrastructure, with malicious objectives. This activity spans energy, finance, transportation and storage, information and communication, and mining and manufacturing sectors. Several cases also involved internet-facing and edge infrastructure exposure, including exploitation of the React2Shell vulnerability, discussed below, and an instance of likely threat actor infrastructure receiving data over HTTP from the firewall of an African energy distributor.

Country-Level Findings, 2024 to 2025

South Africa represented the largest share of **Darktrace-observed** activity in the region, accounting for **45%** of incidents, with network activity making up **67%** of its detections and ransomware appearing in 20% of cases. **Kenya** accounted for **6%** of regional activity, split evenly between email/SaaS and network detections, with ransomware present in **33%** of attacks. Nigeria accounted for a smaller share of overall incident volume, at 5%, but showed the highest concentration of ransomware of the three: present in 80% of investigated cases, alongside data exfiltration indicators in **40%.**

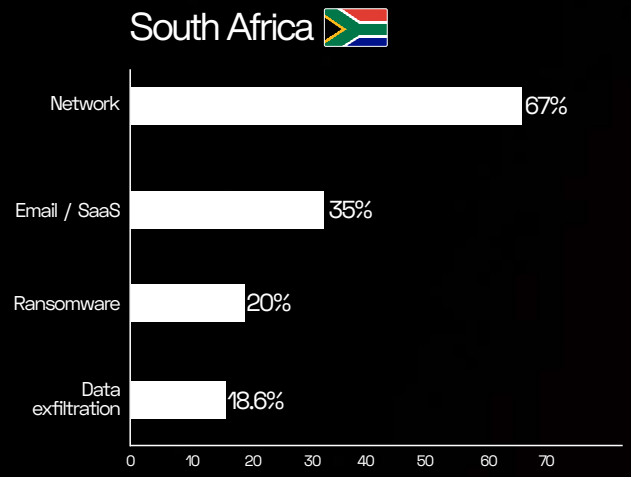

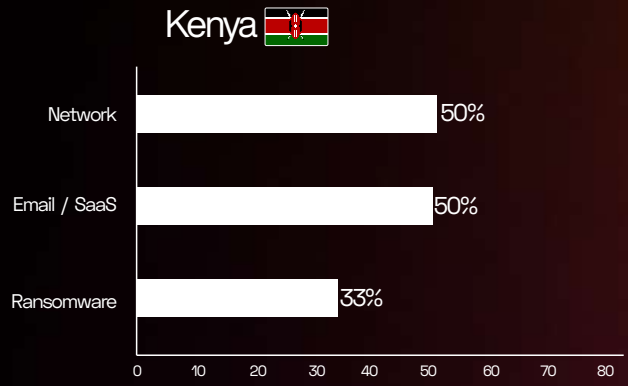

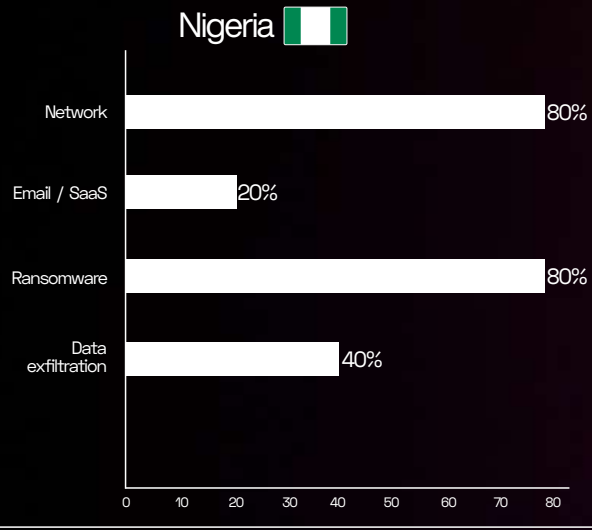

The gap between Nigeria's share of overall incident volume and its ransomware concentration is worth noting. Lower observed volume does not mean lower severity: where Darktrace investigates Nigerian cases, four in five involved ransomware, a materially higher rate than either South Africa or Kenya. This is consistent with the broader pattern raised elsewhere in this report that Nigeria's cyber risk is not proportional to its visibility.

Case Study: React2Shell Across the Region

React2Shell sits directly on the line between the two periods above. The vulnerability was disclosed in December 2025, at the tail end of the regional dataset's window, and its exploitation, including the AI-generated exploit tooling identified later, continued into the H1 2026 period this report covers. It is included here because it shows the regional pattern above, rapid exploitation of internet-facing infrastructure, in motion: not a historical artefact, but the same behaviour carried directly into the current reporting period.

According to Darktrace, multiple organizations across Africa were targeted by attackers exploiting CVE-2025-55812 (React2Shell), a critical vulnerability in React Server Components allowing unauthenticated remote code execution through a single crafted request. Africa accounted for a significant proportion of the React2Shell activity detected globally.

The vulnerability was disclosed and patched on 3 December 2025, but exploitation moved exceptionally quickly. A public proof-of-concept appeared within approximately 30 hours of disclosure, and a honeypot environment run by their research team was compromised within two minutes of deployment, illustrating how rapidly the exploit had been operationalized. Observed attacks primarily targeted internet-facing Next.js applications. Once access was obtained, attackers typically conducted reconnaissance before downloading and executing secondary payloads, including shell-script delivery, HTTP beaconing, and deployment of cryptocurrency miners such as XMRig, with some intrusions showing attempted lateral movement. The campaigns relied heavily on automated scanning for vulnerable hosts at scale.

Cloud-hosted environments were particularly exposed given the prevalence of internet-accessible applications, allowing attackers to move directly from vulnerability exploitation to post-compromise activity without needing phishing or credential theft. Beyond opportunistic crypto mining, the activity extended to financial-sector organizations, with more advanced tooling deployed against higher-value targets. Later cases involved AI and LLM-generated malware and exploitation tooling built specifically to target React2Shell-vulnerable systems, illustrating how generative AI is lowering the barrier to entry for less-skilled operators, a trend directly relevant to the AI-enabled threats discussed elsewhere in this report.

React2Shell is a clear illustration of the shrinking window between vulnerability disclosure and active exploitation, and underscores the importance of rapid patch management, continuous monitoring of internet-facing assets, and behavioural detection capable of identifying malicious activity even where the exploit itself is novel.

Contributors

- Nathaniel Jones - VP Threat Research; Field CISO, AI and Security Strategy Practice
- Emma Foulger - Global Threat Research Operations Lead

Collective Threat Narrative & Cross-Team Findings

When six functions examine the same six months from six different vantage points, patterns emerge that no single team would see alone. Four such patterns define H1 2026. Each is corroborated independently by two or more functions before being stated here as a finding rather than a hypothesis.

I. Nigeria's Risk Is Structural, Not Incidental

Three independent functions converged on the same country this half-year, for three different reasons. Security Testing found Nigerian financial infrastructure under active exploitation, methodical, not sophisticated, exploiting the credential that was reused, and the endpoint left without authorization checks. Darktrace's regional data placed Nigeria at only 5% of Africa's observed incident volume but 80% ransomware concentration within that smaller number, severity entirely disproportionate to visibility. Threat Intelligence supplies the third leg: a named, active data broker advertising Nigerian bank and government data on criminal forums, a national companies-registry agency forced to suspend its own portal, and a formal regulatory advisory landing amid an active probe of 1,369 firms. Three functions, one country, converging from offense, from a global partner's telemetry, and from named threat tracking. That convergence is the report's central finding.

II. The Detection Gap Is Structural, Not Analytical

SOC, DFIR, and Security Engineering independently arrived at the same root cause. DFIR's three major engagements recorded zero automated detections fired, not because the attacks were sophisticated, but because the logs that would have caught them didn't exist: no egress logging, no authentication source-IP retention, Windows Event Logs cleared before anyone arrived. SOC's own assessment names the same gap from the front line: the concern heading into H2 is coverage, not capability. Security Engineering closes the loop from the inside unmaintained EDR and SIEM tooling, and incomplete SIEM source coverage, as the two most consistently observed weaknesses across nine customer engagements. Three functions, one conclusion: the gap is in what gets logged and maintained, not in analyst skill or technology choice.

III. Attackers Are Exploiting Trust, Not Breaking Defences

Five functions now sit on this line, and Threat Intelligence didn't just corroborate it; they named the mechanism directly: trust, not vulnerability exploitation, is now the primary attack surface. TeamPCP's breach of GitHub's own internal repositories is the cleanest illustration available, one poisoned VS Code extension on one developer's machine, no perimeter broken, nothing forced, just a trusted session used exactly as designed. ByteToBreach's own technique profile leads with valid accounts and stolen credentials, not exploitation. Threat Hunting's Trello and Miasma Worm cases relied on legitimate platform infrastructure and a file developer already trusted. Security Testing's most common finding, appearing in 60 to 70% of engagements, was an API returning data because nobody validated what was really being asked for. DFIR's three fraud engagements each ran through something the environment already trusted: a whitelisted IP, a compromised admin account, valid credentials. None of these are exotic techniques. They are trust relationships already granted, used exactly as intended by someone who shouldn't have had them.

IV. AI Is a Two-Way Factor, Now Inside esentry's Own Environments

This is no longer an industry-wide claim borrowed from vendor research. SOC's own analysts observed phishing lures with contextual precision consistent with AI-assisted generation, at the same time as AI-assisted triage became the only way to operate at 3.5 million alerts. Threat Hunting flagged AI-assisted browser extensions as H1's fastest-growing unapproved-installation category.

Threat Intelligence adds the sharpest edge: AI platforms and credentials are now targets, and synthetic identity fraud, a fabricated identity with no real victim to report it, is a fundamentally different detection problem than a stolen one. Darktrace's React2Shell case study shows the same dynamic externally: AI-generated exploit tooling built specifically to target a vulnerability disclosed only months earlier. esentry's own agentic SOC capability, launched this year with human sign-off required at every high-impact decision point, is the organization's direct answer to the defensive half of the same equation.

Telescopic Look into H2 Emerging Trends

AI-enabled cybercrime will keep accelerating.

Automated phishing, malware generation, and reconnaissance are becoming default tooling for attackers of every skill level, not a novelty reserved for the most capable.

Ransomware will increasingly target cloud infrastructure.

An extension of the shift toward data theft and extortion has already been visible throughout H1.

Trusted SaaS platforms will see greater abuse.

The Trello campaign detailed in this report is confirmation that this shift has already arrived, not a prediction that it might.

Automated fraud operations will grow.

including synthetic identity fraud, since a fabricated identity leaves no victim behind to trigger detection.

Software supply-chain attacks will continue.

Automated phishing, malware generation, and The TeamPCP pattern documented in this report, which compromised a trusted developer tool rather than a target directly, is a template other actors will copy, not an isolated event.

Identity-based attacks will keep expanding.

Credential theft has already overtaken traditional malware as the preferred way in. Nothing in H1's data suggests that the trend reverses.

African digital infrastructure will face increased targeting.

Consistent with the disproportionate severity already observed in Nigeria this half-year.

Deepfake-enabled fraud will increase.

against executives, financial transactions, and identity verification processes, a further extension of AI as both a weapon and infrastructure.

Securing the Next Wave

Africa's technology and business ecosystem is expanding rapidly. New ventures, from early-stage companies to fast-scaling enterprises, are digitizing operations, adopting cloud infrastructure, and building customer-facing platforms faster than at any previous point in the region's history. This growth is a genuine economic opportunity. It is also, from a security perspective, where risk accumulates fastest.

Across its broader market engagement, esentry has observed a recurring and concerning pattern: growing businesses account for a disproportionate share of the security weaknesses and compromises we encounter. In many cases, security has been treated as something to be addressed after a product is built, rather than as a requirement that scales alongside the business. The result is that organizations grow quickly in users, revenue, and integrations, while their security posture remains fixed at an early-stage level.

This gap is not merely an internal risk to the organizations themselves. Several of the growing businesses esentry has assessed serve as direct integration points into financial institutions, providing payment processing, data exchange, or platform services that banks and other regulated entities rely on for operations. Where these organizations lack quality security controls, the exposure extends beyond their own systems into the financial institutions connected to them. A weakly secured plug into a bank's ecosystem is, in practice, a weakness in the bank's own attack surface, regardless of where the vulnerable code or infrastructure sits.

The Pattern Typically Looks Like This

- Security as an afterthought — treated as a post-launch concern rather than a design constraint, addressed only once a product is already live and integrated with partners.
- Uneven scaling — growing businesses often scale their user base, transaction volumes and third-party integrations well ahead of scaling their security controls, leaving a widening gap between exposure and defence.
- Inherited institutional risk — where these organizations sit within the financial ecosystem as vendors, partners or integration points, weak controls on their side become an inherited risk for the regulated institutions they connect to.

None of this reflects poor judgement on the part of founders and operators. It reflects the ordinary tension between growth velocity and security maturity that every scaling organization experiences, one that becomes materially more serious when the organization handles financial data or sits inside a bank's technology supply chain.

In response to this pattern, esentry has curated a set of security controls and tooling specifically suited to organizations at this stage of growth, recognizing that the right security investment for a scaling business looks different from what is appropriate for a mature enterprise. As esentry continues to work with organizations at earlier stages of growth, this segment is where we expect to feature more prominently in future reports.

The Eye Acts - Recommendations and Forward Posture - H2 2026

The Eye of Horus does not merely observe. It acts. The following priorities are not drawn from an industry checklist. Every one of them traces something a specific team found in a specific client environment in the past half year. Where the evidence didn't reach far enough to support a recommendation, we left it out rather than pad the list.

Log what you need to prove it happened.

DFIR's number for H1 was not the fraud value; it was the zero next to automated detections fired. Three engagements, three sectors, and not one control caught anything while the attack was live, because the logs that would have resolved the investigation didn't exist. Security Engineering found the same gap from the inside. Mandate egress logging, authentication source-IP retention, and tamper-resistant centralized logging as a baseline, not an upgrade. Then check it. A control nobody maintains is a control that has already failed.

Stop trusting the client to tell the truth about itself.

Security Testing's single most common finding this half-year, appearing in 60 to 70% of engagements, was an API returning more than it should because no one validated what the client claimed to be asking for. Enforce authorization at the backend at every endpoint with the assumption that anything arriving from outside the server is lying until proven otherwise.

Assume the legitimate channel is the attack.

Threat Hunting proved it at one scale: a Trello invitation that passed every email authentication check because it really was Atlassian's mail server, a developer trusting a file already sitting in a repository they controlled. Threat Intelligence proved it on another scale entirely: a poisoned VS Code extension gave TeamPCP access to 3,800 of GitHub's own internal repositories. Vet developer tooling and IDE extensions with the same discipline you'd apply to production access, because increasingly, that's exactly what it is.

Segment the inside as seriously as you guard the perimeter.

Security Testing's internal engagements told a different story from the external ones. Once a foothold existed, weak segmentation and Active Directory hygiene did the rest of the attacker's work for them.

Don't outsource your visibility to someone else's promise.

One of DFIR's three fraud cases ran through a payment provider that processed hundreds of fraudulent transactions with no velocity alert and proactive notice. If a third party sits inside your trust boundary, verify what they will catch. Don't assume it.

Patch on the attacker's clock, not yours.

SOC watched exploitation land within 72 hours of disclosure. Darktrace watched a proof-of-concept appear within 30 hours globally. ByteToBreach's own pattern is opportunistic exploitation of unpatched, internet-facing platforms. Three independent sources, the same lesson: time-to-patch is a survival metric now, not a maintenance one.

Govern the agent the way you would govern a new hire with admin rights. esentry's own agentic SOC

capability launched this year with human sign-off required at every high-impact decision point, and that commitment came before we recommended anyone else make it. Any organization granting an AI system standing permission over money, infrastructure, or high-stakes action should audit that permission chain specifically.

Track named threats before they become the next regulatory bulletin.

One of DFIR's three fraud cases ran through a payment provider that processed hundreds of fraudulent transactions with no velocity alert and proactive notice. If a third party sits inside your trust boundary, verify what they will catch. Don't assume it.

EYE OF HORUS — H1 2026

The Eye Never Closes.

Six teams, one threat landscape, seen from six angles that, on their own, would each have told an incomplete story. Put together, H1 2026 resolves into a single pattern: attackers are not breaking in; they are being let in through a trusted developer tool, a legitimate SaaS invitation, a credential nobody rotated, a log that was never enabled. Nigeria sat at the center of that pattern this half-year, not by coincidence, but because the gap between exposure and visibility is wider here than the numbers alone suggest.

A single new threat will not define H2 2026. It will be determined by whether the trust that enabled H1's intrusions is examined or extended by default again. esentry's answer, across every team in this report, is the same one it gave to its own use of AI this year: build the capability, then govern it before you trust it.

The Eye does not close because the work does not end. Neither does ours.

— esentry

Contact Us

services@esentry.io

212/214 Herbert Macaulay Way.
Yaba, Lagos, Nigeria