2026 Global Threat Landscape Report

Insights from FortiGuard Labs

FORTINET

2026 Global Threat Landscape Report

A Report by FortiGuard Labs

Arturo Torres: Director, FortiGuard LATAM
Douglas Santos: Director, Advanced Threat Intelligence
Derek Manky: VP, Global Threat Intelligence

Collaborators

Mark Robson: Principal Threat Analyst (IR Team)
Ankit Chauhan: Lead Cyberthreat Intelligence Analyst R&D (FortiRecon Team)
Christopher Hall: Principal Cloud Security Researcher (FortiCNAPP Team)
Vijay Dontharaju: Director, Security Engineering (FortiNDR Cloud Team)
Motti Elloul: Director of Product Management & Incident Response, FortiMail Workspace Security

Table of Contents

2026 Global Threat Landscape Report

This report is derived exclusively from **FortiGuard Labs threat intelligence**, leveraging telemetry from millions of sensors deployed worldwide since 2002. It covers data gathered in 2025 (or the most recent 12-month window available per dataset) across multiple security domains and vectors of compromise. Each insight includes a Source Tag that indicates the telemetry origin and is mapped to MITRE ATT&CK to ensure a defensible, repeatable analytical baseline. Findings and recommendations are prioritized by probability and prevalence in observed activity, with a direct focus on detection, response, and automation outcomes for SOC, DFIR, and CISO audiences.

In 2026, the threat landscape cannot be accurately described through isolated indicators or single-domain trends. Adversaries operate across an end-to-end lifecycle that begins well before intrusion through exposure discovery, access brokerage, and industrialized preparation, and continues through exploitation, persistence, monetization, and operational impact.

To reflect this reality, the *2026 Global Threat Landscape Report* introduces the FortiGuard SecOps Kill Chain, a telemetry-driven model built on a foundational advantage. Fortinet delivers SecOps technologies across the attack lifecycle, generating real-world visibility across multiple vectors of compromise. This unified, multi-domain telemetry enables a defensible threat narrative grounded in evidence. The framework provides a consistent analytical structure anchored in MITRE ATT&CK as a common language, while translating telemetry into decisions aligned with continuous threat exposure management (CTEM).

FortiGuard Security Operations (SecOps) chain phases and how to read them

The model describes threat activity across six repeatable stages:

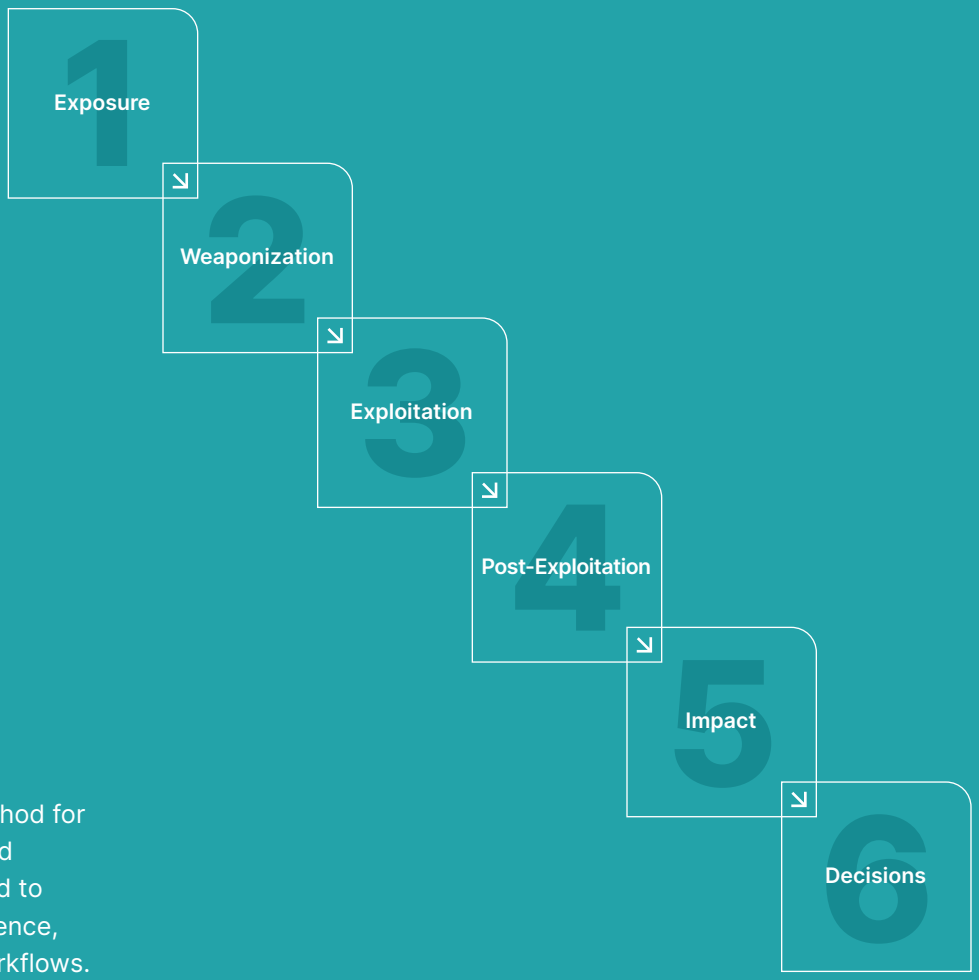

This structure provides a consistent method for correlating telemetry across domains and ensures that each insight can be mapped to ATT&CK techniques, validated with evidence, and operationalized through SecOps workflows.

Introducing FortiGuard SecOps action boxes

To ensure the report delivers operational outcomes, not just insight, every phase of the FortiGuard SecOps Kill Chain is anchored by a standardized FortiGuard SecOps action box. The action box is the primary execution artifact of the report and represents the point where telemetry-driven intelligence is converted into concrete, role-specific action.

The action box is intentionally visual, modular, and reusable. It is designed to be consumed independently of the surrounding narrative and to be directly repurposed into:

- SOC playbooks
- DFIR investigation checklists
- CISO briefings and CTEM planning

To reduce repetition across the report, we reference the integrated FortiSOC stack as the execution layer behind the SecOps action boxes.

FortiSOC combines three operational capabilities:

1

FortiAnalyzer:
log analytics, baselining, reporting, and investigation pivots

+

2

FortiSIEM:
multi-source correlation, rules, incident/case management

+

3

FortiSOAR:
orchestration, playbooks, enrichment, and automated response

NOTE: While FortiSOC is used in the action boxes to improve readability, customers can leverage either the unified FortiSOC SaaS offering or individual FortiAnalyzer, FortiSIEM, and FortiSOAR components.

Prevention through Disruption: Breaking the Industrial Cybercrime Supply Chain

Modern cybercrime now operates as an industrial system. It relies on shared infrastructure, reusable tooling, established marketplaces, and persistent coordination across criminal ecosystems. In that environment, prevention is no longer limited to blocking individual attacks at the point of execution. It also means disrupting the upstream systems that make large-scale cybercrime possible in the first place.

Fortinet's approach to prevention reflects this reality. In addition to detection and response across enterprise environments, Fortinet works directly with law enforcement agencies, industry partners, and global coalitions to identify, map, and dismantle the infrastructure and services that sustain cybercrime operations. This work targets the cybercrime supply chain itself: command-and-control networks, fraud infrastructure, credential theft ecosystems, scam operations, and the marketplaces that connect them.

Over the past year, Fortinet has played an active role in numerous coordinated international disruption efforts, including **INTERPOL Serengeti 2.0** and **Operation Red Card 2.0**, the **Cybercrime Atlas** initiative with the **World Economic Forum**, working with cybersecurity peers through the **Cyber Threat Alliance** (CTA), and a new **Cybercrime Bounty program** launched in partnership with Crime Stoppers International. These initiatives focus on translating threat intelligence into operational action by identifying criminal infrastructure, supporting coordinated takedowns, and reducing the availability of services that enable large-scale abuse.

The **Cybercrime Bounty** program supports this same goal by expanding visibility into active criminal infrastructure and tooling.

Operations such as Red Card 2.0 and Serengeti 2.0 demonstrate how intelligence-driven coordination can disrupt cybercrime at scale. These efforts led to arrests, infrastructure seizures, and the dismantling of fraud and cybercrime operations that depend on shared hosting, communication platforms, and monetization pipelines. From a defensive perspective, this kind of disruption removes entire classes of threats before they ever reach enterprise networks.

Disruption also serves a broader purpose. By pushing investigations toward real operators and real infrastructure, these efforts increase accountability and introduce friction into criminal ecosystems that depend on anonymity and reliability. When infrastructure is repeatedly dismantled and operators are exposed, the cost and risk of operating at scale increases. Over time, this creates deterrence effects by making large-scale cybercrime less predictable, less stable, and less profitable.

The Cybercrime Atlas and CTA extend this model by improving how threat intelligence is shared, correlated, and operationalized across the industry. Their purpose is not only faster detection, but faster disruption. By mapping adversary infrastructure, relationships, and operational dependencies, these efforts enable the identification of systemic weaknesses in the cybercrime ecosystem and coordinate action against them.

The Cybercrime Bounty program supports this same goal by expanding visibility into active criminal infrastructure and tooling. It helps surface high-impact targets for investigation and takedown and shortens the time between discovery and action.

Taken together, these efforts represent a shift in how prevention must be understood. In an industrialized threat landscape, prevention is not only about stopping individual exploits or blocking specific malware families. It is about reducing the capacity of the adversary ecosystem itself by increasing costs, adding friction, and degrading the reliability of the shared services that enable scale.

This upstream disruption work complements the technical controls described throughout this report. While exposure management, detection, and response remain essential, the most effective form of prevention is to ensure that fewer threats ever reach operational maturity in the first place. When criminal infrastructure is dismantled, marketplaces are disrupted, and operators are forced to adapt, and the entire attack pipeline becomes less reliable.

This is a critical control point in the attack chain. While implementing cyber defenses here does not eliminate the need for downstream defenses, it significantly reduces the volume, velocity, and efficiency of the threats those defenses must handle.

Executive Summary

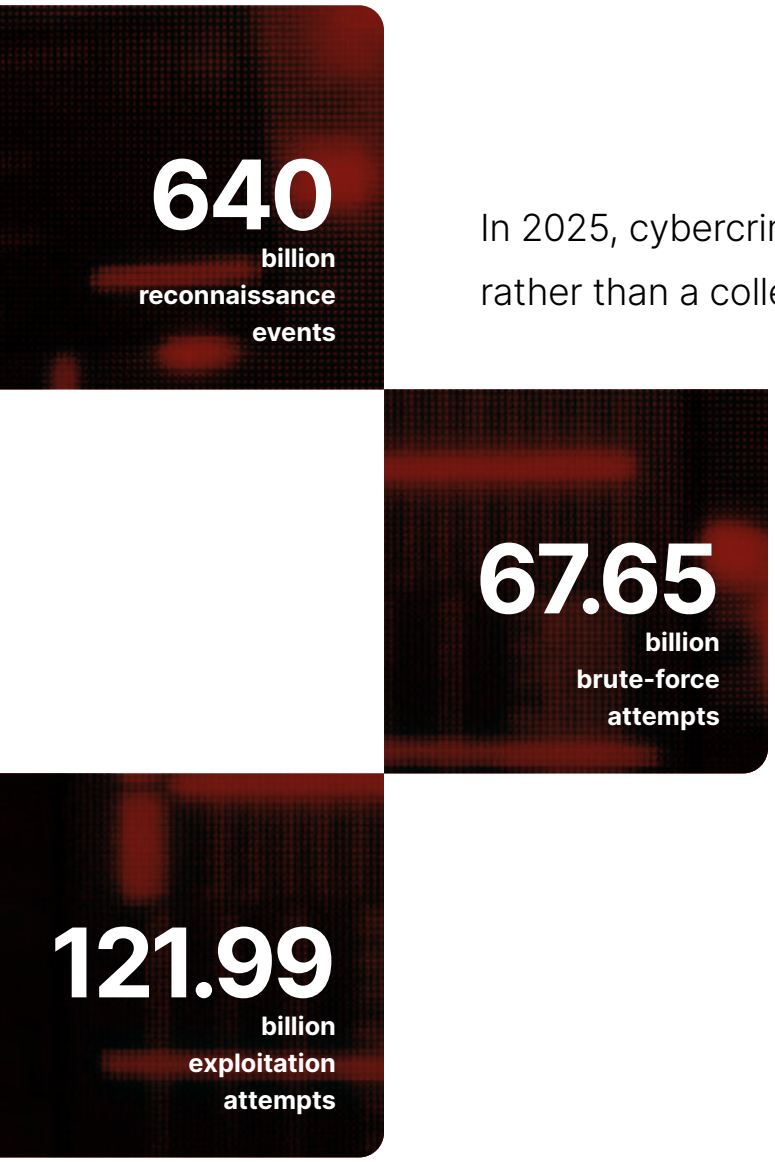

In 2025, cybercrime operated as an industrial system rather than a collection of isolated campaigns.

Fortinet telemetry recorded **640 billion reconnaissance events (YoY -45%), 67.65 billion brute-force attempts (YoY -22%)**, and **121.99 billion exploitation attempts (YoY +25%)** globally. For context, exploitation attempts were just over **97 billion in 2024**, indicating a substantial year-over-year increase in automated, at-scale attacks. This reflects a continuous, machine-speed operating model in which exposure is persistently mapped, validated, and activated without reliance on campaign cycles or manual intervention.

Exploitation is increasingly driven by **availability and automation rather than new exploit development**. Of the 635 vulnerabilities observed under active exploitation, **53.86% had publicly available proof-of-concept code and 31.18% had fully functional exploit code**. These ratios are broadly consistent with 2024 levels and show how quickly attackers operationalize existing exploit material once it becomes available.

Time-to-exploit continues to shrink as automation accelerates this process. Where exploitation activity in earlier reporting periods often appeared **around a week after disclosure**, in 2025 TTE was consistently observed to occur **within 24–48 hours**, consistently outpacing traditional patch and remediation timelines.

Decision Ask:
Organizations should adopt CTEM approaches to better understand their attack surface and prioritize remediation. In an increasingly **industrialized threat environment**, defensive velocity—measured through time to detect, time to contain, and time to revoke compromised credentials—should be treated as a primary business risk indicator. Security advantage is no longer defined by the number of tools deployed, but by how quickly organizations can identify and stop intrusions.

Identity exposure remains the upstream fuel of industrialized intrusion. FortiRecon observed **4.62 billion stealer logs** traded or shared on the darknet, **a 79.07% increase compared to 2024**, with approximately **1.7 billion** stolen credential records circulating in 2024, indicating continued expansion of the credential economy. In cloud environments, identity compromise remains the dominant intrusion vector, with valid credentials serving as the exploit and APIs as the execution layer.

Ransomware continues to operate as a mature production model. **7,831 confirmed victims** were recorded globally in 2025, compared to approximately **1,600** in 2024 **(389% YoY)**, with activity sustained throughout the year. Impact is systematic, economically optimized, and repeatable rather than episodic.

Execution patterns reinforce the same industrial logic. EDR telemetry shows that **48.96%** of suspicious activity is tied to the abuse of legitimate applications (LOLbins), confirming that scale is achieved through native tooling and automation rather than bespoke malware. Post-compromise control is sustained through persistent infrastructure: **7.10 billion botnet C2 detections** (approximately **19.4 million per day) (YoY -11%)**. This reflects an operational command layer that maintains access and enables staging, persistence, and monetization.

Across exposure, exploitation, execution, and impact, the common variable is not just attacker sophistication. It is speed and reuse.

Across exposure, exploitation, execution, and impact, the common variable is not just attacker sophistication. **It is speed and reuse.**

Key findings: impact-driven, numbers-first

- **Exposure: cybercrime industrializes opportunity**
  640B reconnaissance events and 67.65B brute-force attempts confirm continuous, machine-speed mapping and credential conversion at global scale.

- **Weaponization: identity and exploits as industrial inventory**
  4.62B stealer logs and 635 actively exploited vulnerabilities (53.86% with public PoC) show exploit packaging and credential commoditization at scale.

- **Exploitation: execution as a race condition**
  In 2025, 121.99B exploitation attempts and 57.32% of vulnerabilities that entered exploitation show that readiness and automation, not CVSS, define impact. **TTE has compressed from a few days to zero**, with multiple critical outbreaks reaching first exploitation signals the **same or the next day**.

- **Post-exploitation: persistent command infrastructure**
  7.10B botnet C2 detections (~19.4M/day) confirm compromise is sustained through industrialized command layers, not isolated access.

- **Impact: ransomware as high-throughput production**
  7,831 confirmed victims in 2025, concentrated among scalable groups, demonstrate that ransomware is a continuous economic engine, not an episodic campaign.

- **Cross-threat convergence: one vulnerability, multiple outcomes**
  The same vulnerabilities are reused across threat activity: 22.83% are leveraged only in ransomware campaigns, 19.53% only in APT operations, and 20.47% in both, while 43.15% are exploited at scale. As these categories are not mutually exclusive, this convergence highlights how individual vulnerabilities can simultaneously enable extortion, espionage, and large-scale compromise.

- **Execution model: fileless, native, automation-ready**
  48.96% LOLbin abuse, ~8% injection/hollowing, and 11.5% immediate sensitive data access confirmed via native tooling, not bespoke malware.

- **Cloud: identity as the exploit, APIs as the engine**
  Identity-driven compromise, discovery-heavy API bursts, and monetization via AWS SES abuse, cryptomining, and resource hijacking define industrialized cloud intrusion.

In 2025, cybercrime ran at an industrial scale.

635 vulnerabilities were exploited.
7,831 organizations were extorted.

Automation, identity abuse, and patch latency, not zero days, defined impact.

As AI accelerates reconnaissance, weaponization, and execution, **TTE is now collapsing to same-day windows, often within 24 hours.**

Defensive velocity is no longer a technical metric. It is a business decision.

Exposure as an Industrial Input: How Cybercrime Industrializes Opportunity

In an industrialized threat landscape, exposure is no longer an accidental weakness. It is a resource that adversaries actively harvest, refine, and operationalize at scale. Credentials, access paths, vulnerable services, and misconfigurations are collected, validated, packaged, and reused through a mature supply chain that feeds continuous intrusion activity across the attack lifecycle.

This shifts the defender's problem. Exposure is no longer defined by what is visible today, but by what has already been captured, cataloged, and redistributed through underground markets and automated tooling. Identity data, access credentials, and infrastructure intelligence persist long after the original compromise, allowing attackers to bypass discovery and move directly to execution. The risk is no longer just what defenders fail to hide. It is also what adversaries already possess and can deploy on demand, at machine speed and industrial scale.

Darknet landscape: exposure already harvested (FortiRecon intelligence)

In an industrialized threat landscape, exposure no longer begins with reconnaissance. It begins in underground markets. FortiRecon intelligence confirms that a significant portion of modern cyber exposure is pre-harvested, structured, and commercialized through darknet IABs ecosystems before any new scanning or brute-force activity takes place. (In the context of cybersecurity, IABs are threat actors who specialize in gaining unauthorized access to corporate networks and then selling that "entry" to other cybercriminals—most notably ransomware operators.) Databases, credentials, validated access paths, and attacker tooling are continuously advertised and exchanged, forming an upstream supply chain that feeds downstream intrusion activity. As this supply chain becomes more automated and AI-assisted, the time between data theft, packaging, and operational use continues to shrink, directly compressing initial access across multiple intrusion paths.

Identity and data as pre-existing exposure

Darknet database offerings in 2025 were dominated by datasets optimized for **reuse and automation**, not one-off exploitation:

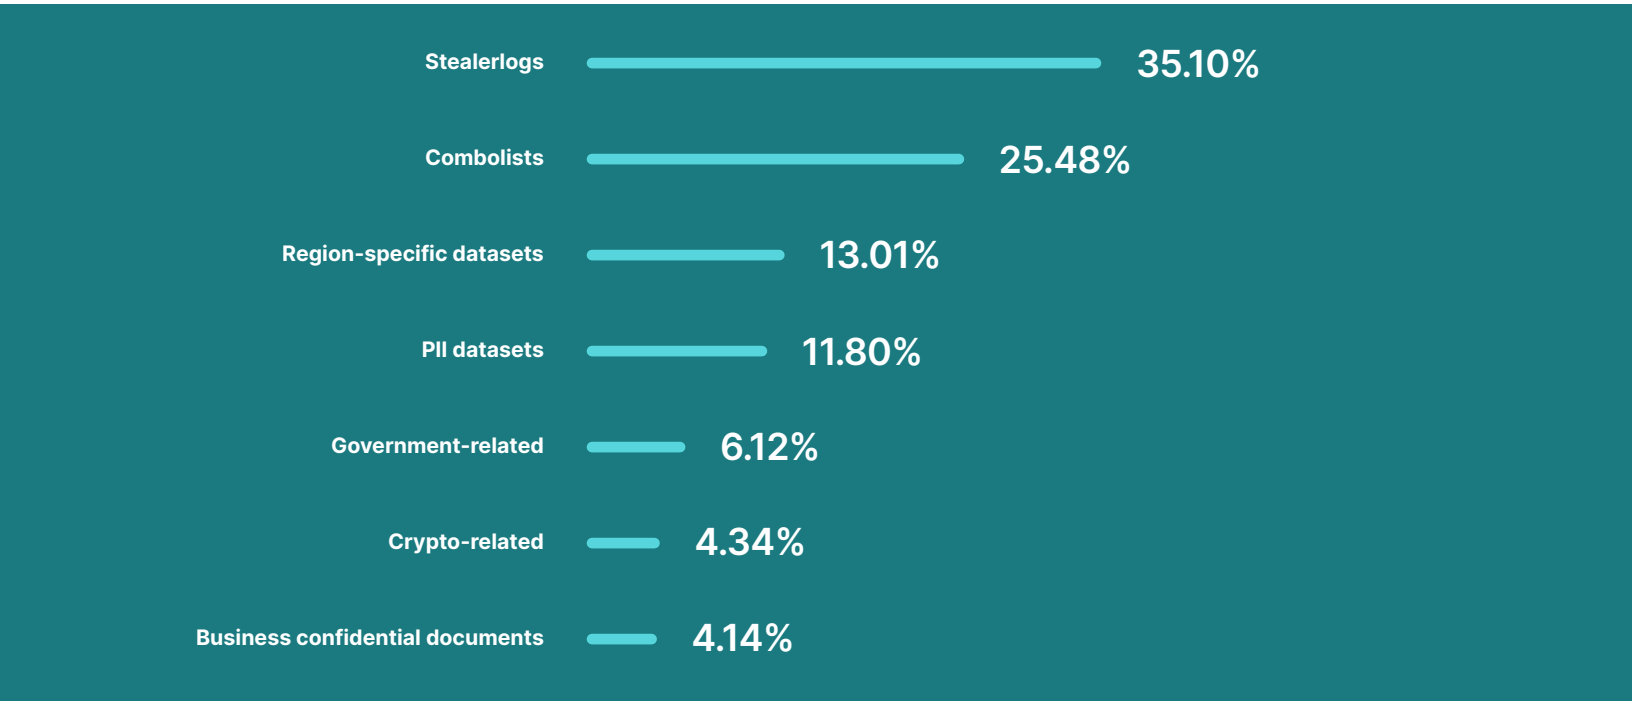

This distribution reflects a mature criminal supply chain. Stealer logs and combolists serve as high-volume fuel for credential reuse, enabling credential stuffing and account takeover attempts across VPNs, SSO portals, webmail, and other remote services, often without any additional discovery effort.

Trading activity was highly concentrated, led by threat actors **JHON02 (13.54%**) and **BESTCOMBO (12.84%)**, followed by **MEGACLOUD / CODER / UNIQUECOMBO / VALIDMAIL**, confirming that exposure harvesting operates as a production market rather than an opportunistic ecosystem.

In this model, **identity is not a target. It is raw material**. Increasingly, AI-assisted tooling accelerates validation, sorting, and reuse of that material, further reducing initial access time once credentials enter circulation.

Access brokerage: exposure converted into entry points

Beyond data, IABs actively trade validated access into enterprise environments. The most frequently advertised access types are:

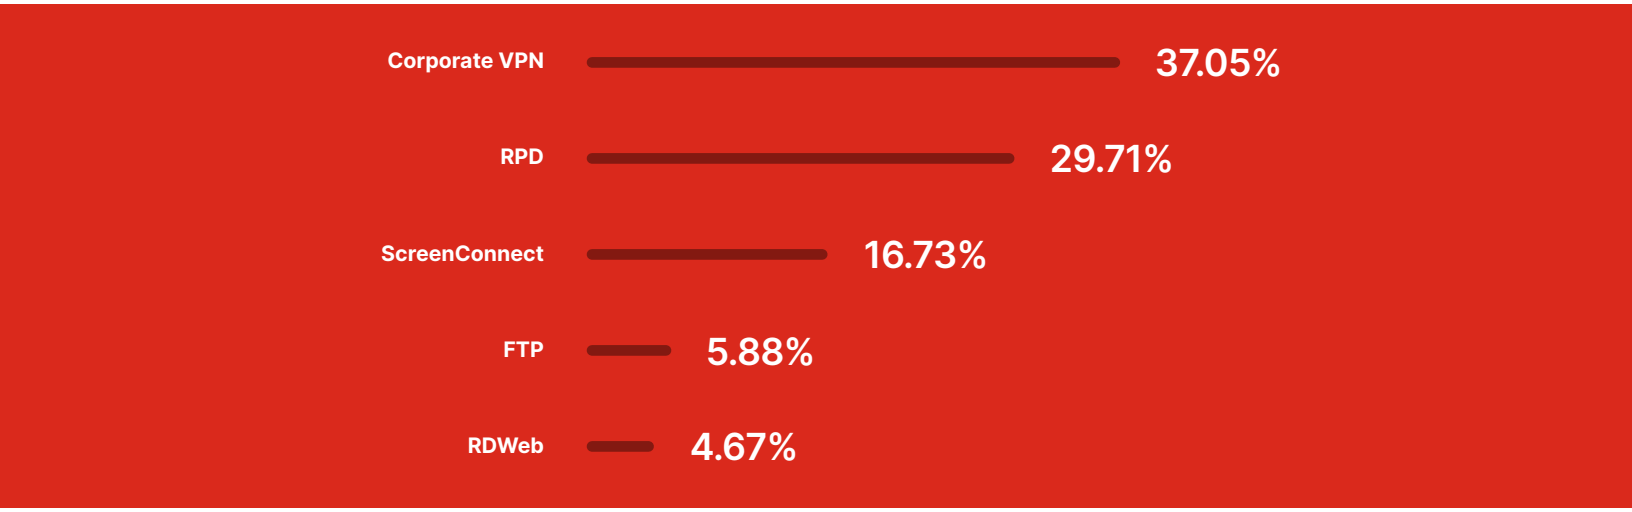

Additional categories, including **SSO (0.96%), AnyDesk (0.86%), Email access (0.70%), Cloud access (0.59%), SSH (0.51%)**, and smaller volumes of webshell, API, admin panel, and storage access, demonstrate the breadth of pre-positioned intrusion options available.

Access advertising is dominated by a small set of sellers, with **WWW (30.10%)** leading observed activity (followed by **XRAHITEL / TALISMANUDACHIZZ / SHELLSHOP**), reflecting an access-broker economy optimized for speed, scale, and reliability.

In an industrialized model, intrusion does not begin with exploitation. It begins where access is already established.

Credential stealers: industrial-scale identity production

Credential-stealer malware remains a primary upstream engine for exposure generation. FortiRecon telemetry shows stealer activity dominated by:

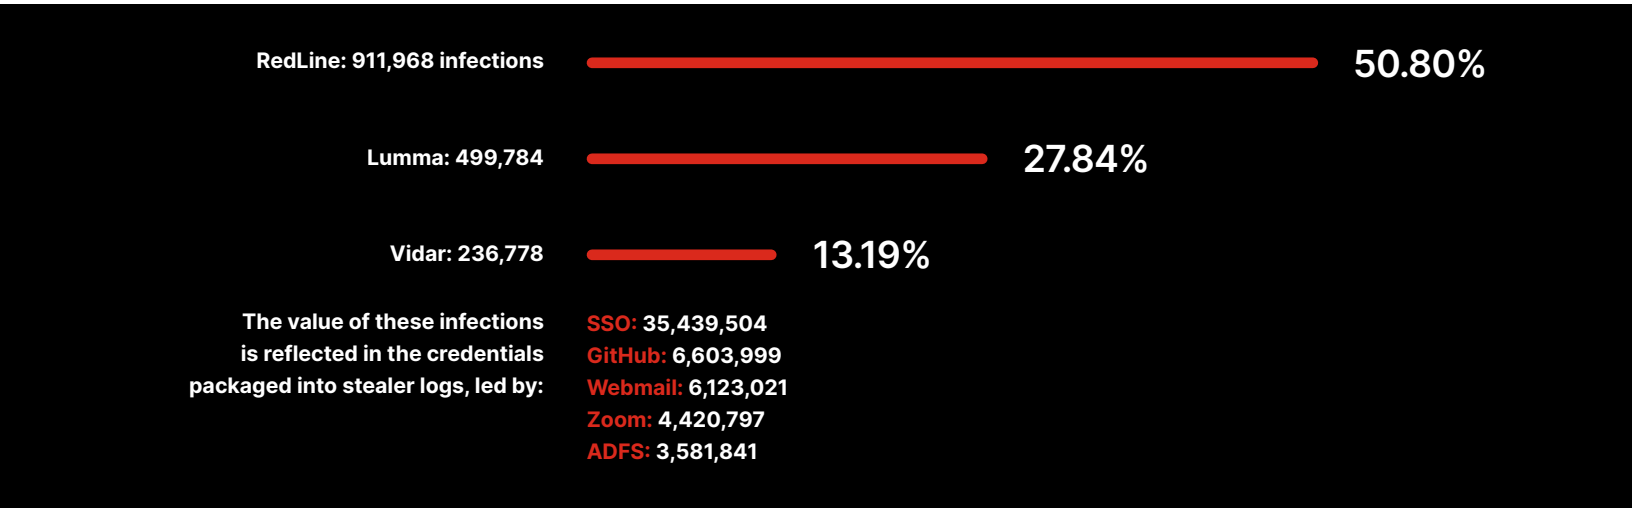

Additional enterprise platforms appear consistently across stealer datasets, spanning **Business/ERP (Oracle, SAP Portal), Dev/Repo (GitLab), IT/Collaboration (Jira, Confluence), Auth/Identity (Auth Portals, Okta, ADFS), Remote Access (AnyDesk, Citrix), Email (OWA, Zimbra), SaaS (Salesforce),** and **Content (SharePoint)**.

This confirms a structural shift: **identity exposure is no longer episodic**. It is **continuously produced, packaged, and redistributed** through the infostealer supply chain. As AI-assisted phishing, malware customization, and log processing mature, this production cycle continues to shorten, further compressing TTE across identity-driven intrusion paths.

Cyber reconnaissance at machine scale: automated scanning as the foundation of industrialized attacks (IPS intelligence)

Active scanning remains a dominant entry point in the modern Cyber Kill Chain. Throughout 2025, FortiGuard Labs telemetry confirmed that automated reconnaissance continued to operate at sustained, planetary scale, forming the foundation of nearly all downstream weaponization and exploitation activity.

While total scanning volumes declined year over year, this does not indicate reduced adversary capability or intent. It reflects operational optimization. Threat actors are scanning more selectively and efficiently, with clearer downstream objectives, and rely on persistent visibility into exposed attack surfaces rather than indiscriminate probing.

Across all regions, observed scanning activity aligns predominantly with **MITRE ATT&CK TA0043 (Reconnaissance)**, with sustained concentration in:

- **Active Scanning (T1595)**
- **IP Block Scanning (T1595.001)**
- **Vulnerability Scanning (T1595.002)**
- **Victim Network Information Gathering (T1590)**, including **DNS Enumeration (T1590.002)**

Executed continuously and at machine scale, these techniques form an industrialized reconnaissance layer rather than isolated probing. This layer enables attackers to systematically map exposed infrastructure, refresh target intelligence, and maintain a persistent pipeline of opportunity that feeds credential abuse, weaponization, and the exploitation of workflows.

As AI-assisted targeting, classification, and prioritization mature, this reconnaissance layer becomes more tightly coupled to exploit packaging and access brokerage, further reducing TTE from discovery to impact.

This confirms a critical shift. **Cyber reconnaissance is no longer a preparatory phase. It is a continuous, automated capability tightly coupled to adversary weaponization pipelines.**

Global reconnaissance volume in 2025

FortiGuard telemetry observed approximately **640 billion scanning events** over the year, roughly **1.75 billion per day, 72–75 million per hour**, and **more than 20,000 per second** sustained (-45% YoY). At this velocity, adversaries do not need to initiate scanning when new vulnerabilities are disclosed. The global attack surface is already mapped, continuously refreshed, and maintained in an operational readiness state.

This model compresses the time between disclosure and exploitation. When exploit code becomes available, targeting can begin immediately without a preparatory discovery phase. Reconnaissance is no longer a discrete step in the attack chain. It is a continuous background process that feeds weaponization and exploitation pipelines in near real time, placing defenders at a structural disadvantage if exposure is not proactively managed.

Regional volumes reflect the same execution model at different scales: **APAC (~260.13B)**, **EMEA (~168.36B), North America (~156.68B)**, and **LATAM (~54.65B)** scanning detections in 2025. While absolute volumes vary, the operational pattern is globally consistent—continuous scanning, protocol prioritization, and systematic reuse of reconnaissance data to accelerate exploitation.

The result is a persistent reconnaissance layer that shortens attacker decision cycles and further compresses the disclosure-to-exploitation window. In an industrialized threat environment, defenders are no longer racing individual campaigns. They are competing against an always-on, machine-scale mapping operation that continuously collects, validates, and updates exposure intelligence.

Region	Annual scanning detections (2025)	Per day	Per hour	Per second	Operational dynamic
APAC	260.13B	~713M	~29.7M	~8,250	Pronounced end-of-year acceleration, particularly against SIP services and DNS infrastructure, consistent with dense internet-facing services, telecom infrastructure, and cloud adoption.
EMEA	~168.36B	~461M	~19.2M	~5,300	Lower overall noise but higher consistency, dominated by DNS enumeration and vulnerability-focused probes recon feeding precision exploitation workflows.
North America	~156.68B	~429M	~17.9M	~5,000	Lower volume but higher targeting efficiency; reconnaissance aligned to vulnerability discovery, DNS mapping, and public-facing service enumeration, adaptation to mature defensive environments.
LATAM	~54.65B	~150M	~6.2M	~1,730	Largest YoY reduction, yet persistent and highly focused on SIP, DNS, and repeatedly exposed assets, suggesting reuse of historical reconnaissance data rather than reduced adversary interest.

> Across all regions, threat actors are not randomly probing the internet. They are systematically feeding an industrial reconnaissance pipeline."

Reconnaissance targets and automation: what cybercrime scans for and how it industrializes discovery

Across all regions, threat actors are not randomly probing the internet. They are systematically feeding an industrial reconnaissance pipeline. FortiGuard telemetry shows consistent prioritization of infrastructure categories that maximize scale, reuse, and monetization, reflecting a mature, production-oriented cybercrime model. In this model, reconnaissance is not exploratory. It is inventory creation, where exposed services, access points, and protocols are continuously identified, indexed, and refreshed so they can be reused and activated on demand. This compresses the time between vulnerability disclosure and exploitation, tying discovery directly to downstream weaponization and exploitation workflows.

Target (what they scan)	Description (per telemetry / narrative)	Downstream impact (what it enables)	Defender impact (why it matters)
SIP / VoIP exposed services	Exposed VoIP services with weak authentication controls and persistent exposure.	Call fraud, lateral movement, and resale within access-broker ecosystems.	Increases the probability of repeatable entry points; forces SIP exposure to be treated as a continuous operational risk.
DNS exposed infrastructure	DNS infrastructure continuously scanned to map environments and identify exploitable domains.	Preparation for C2 and traffic tunneling; mapping/selection of exploitable domains; multi-campaign reuse.	Compresses reaction windows: adversaries arrive with domain/infrastructure context; increases the value of DNS hardening and monitoring.
Public-facing services & management interfaces	Internet-facing services and admin interfaces scanned to inventory exploitable conditions.	Inventories of vulnerable versions, misconfigurations, and exposed control planes activated on-demand during weaponization.	Converts exposure into "ready inventory;" without continuous exposure management, defenders react late (after activation begins).

Threat actors no longer rely on ad hoc or manually driven scanning. In 2025, FortiGuard Labs telemetry showed automated reconnaissance tooling had become a force multiplier, enabling cybercrime to industrialize target discovery and sustain continuous visibility into exposed infrastructure. Rather than developing bespoke tools, attackers overwhelmingly favor commodity and dual-use scanning frameworks orchestrated through automation to optimize speed, coverage, and reuse. This decouples reconnaissance from exploitation, allowing scanning results to persist as actionable intelligence that can be weaponized on demand.

Individually, these tools are widely available and well understood. Their strategic impact emerges from how they are combined, automated, and reused. In 2025, reconnaissance tooling functioned less as a set of discrete utilities and more as infrastructure within the cybercrime economy, continuously feeding access brokerage, exploitation frameworks, and monetization workflows. The result is that modern reconnaissance is not an event to react to. It is a persistent condition defenders must manage.

Tools consistently dominating global reconnaissance telemetry during 2025 included:

- **Commercial vulnerability scanners (such as Qualys)**, reflecting the dual-use nature of enterprise tooling leveraged at scale to identify vulnerable software versions and misconfigurations, and reduce the effort required to prepare exploitation pipelines.
- **Nmap**, a foundational component for large-scale port scanning, service fingerprinting, and protocol identification, is particularly useful when executed continuously rather than episodically.
- **Nessus and OpenVAS**, observed at lower volumes but used to enrich reconnaissance data with vulnerability context that supports prioritization and downstream weaponization.

> At this scale, brute force is not background noise. It is industrial infrastructure."

Brute force as a production pipeline at industrial scale

If reconnaissance identifies an opportunity, brute-force industrializes access. In 2025, FortiGuard Labs telemetry confirmed that brute-force attacks remained one of the most scalable, reliable, and cost-efficient mechanisms for converting exposed services into valid access. What differentiates modern brute force is not the technique itself, but how it has been industrialized—automated, optimized, and executed continuously at a global scale. Brute force no longer operates as a noisy, opportunistic tactic. It functions as a production process designed to maximize credential conversion while minimizing attacker effort.

Across all regions, observed activity aligns overwhelmingly with **MITRE ATT&CK TA0043 (Reconnaissance)**, with sustained concentration in **Active Scanning (T1595)**, **IP Block Scanning (T1595.001)**, **Vulnerability Scanning (T1595.002)**, **and Victim Network Information Gathering (T1590)**, including **DNS Enumeration (T1590.002)**.

Executed continuously and at internet scale, these techniques form a persistent, automated reconnaissance layer that allows attackers to systematically identify, classify, and refresh exposed attack surfaces, directly feeding downstream credential abuse, weaponization, and exploitation workflows. These techniques appear repeatedly immediately after reconnaissance, forming a repeatable and reliable access pattern rather than isolated attempts.

During 2025, FortiGate IPS telemetry recorded approximately **67.65 billion brute-force events globally (YoY -22.1%)**. This activity translates into:

- **~185 million attempts per day**
- **~1.3 billion attempts per week**
- **~5.6 billion attempts per month**

As with reconnaissance, the year-over-year decline reflects efficiency gains rather than reduced threat activity. Threat actors are making fewer attempts against better-selected targets, increasing the probability of success per credential tested.

At this scale, brute force is not background noise. It is industrial infrastructure.

Brute-force activity shows a concentrated regional distribution, reflecting where exposed credential-based services and reuse patterns provide the highest return:

- **EMEA (~40.9%) → ~27.7B annually (~76M/day)**
- **APAC (~38.3%) → ~25.9B annually (~71M/day)**
- **North America (~15.5%) → ~10.5B annually (~29M/day)**
- **LATAM (~5.3%) → ~3.6B annually (~10M/day)**

Despite differences in absolute volume, the execution model is consistent across regions, where sustained brute-force pressure is applied to services with predictable credential exposure.

Services under sustained brute-force pressure

Telemetry shows that brute-force operations overwhelmingly target credential-centric services that reliably convert attempts into access:

- **SMB: 25.1B** events **(37.9%)**
- **SSH: 14.8B** events **(22.4%)**
- **MySQL: 7.3B** events **(11.1%)**
- **RDP: 6.7B** events **(10.1%)**

These services offer the shortest operational path to initial access because of their ubiquity, frequent credential reuse, and the immediate access they provide once compromised. From an industrial perspective, they function as high-throughput access points. They are easy to automate, inexpensive to attack, and consistently exploitable.

Why brute force converts so reliably

Brute-force success does not rely on sophistication. It relies on scale and probability. Attackers exploit:

- **Credential reuse across exposed services**
- **Weak or missing rate limiting**
- **Inconsistent enforcement of MFA**
- **Predictable authentication workflows**

By applying sustained pressure at machine speed, attackers convert statistical likelihood into reliable access over time, particularly in environments where identity controls are weak or unevenly deployed.

FortiGuard SecOps action box: exposure

In an industrialized threat landscape, exposure is harvested before exploitation begins. When reconnaissance and brute force are combined, defenders still have the advantage, but only if they act before access is gained.

SOC: Signals · Correlations · Detections · Automation

Phase	SecOps-Ready Actions
Signals to Monitor	▪ Enrichment-focused reconnaissance (DNS enumeration, SIP scanning, service fingerprinting): flag spikes vs baseline on internet-facing services [Detection: FortiGate IPS] ▪ Sustained brute-force attempts against credential-based services (SMB, SSH, RDP, MySQL): track failure bursts/sprays per service/host [Detection: FortiGate] ▪ Reuse of infrastructure (ASN, IP clusters, tooling fingerprints): cluster repeat sources across targets to raise confidence. Optional confidence boost: external infrastructure/exposure context. [Enrichment: FortiRecon]
Correlation Logic	▪ Reconnaissance followed by brute force against the same service or asset class: chaining rule within a short window. ▪ Temporal proximity between scanning bursts and authentication abuse: burst→spike correlation using baselines. ▪ Repetition of the same access pattern across multiple targets: group spray/distributed patterns. [Correlation/Case execution: FortiSOC]
Detection Gating (Operational)	▪ Low: isolated scanning or brute force (no chaining) ▪ Medium: recon + brute force (no reuse/repeat) ▪ High: recon→brute force chaining + reuse or repeated targeting
Automation & Response	▪ SOAR enrichment linking recon source → targeted service → asset criticality/owner: enrich then route. ▪ Correlate brute-force attempts → identity class → privilege potential: prioritize admin/service identities. ▪ Automated controls when High is met: contain (rate-limit/restrict) then harden credentials. [Orchestration: FortiSOC]

DFIR: Evidence · Artifacts · Reconstruction (compact)

Phase	SecOps-Ready Actions
Evidence to Collect	▪ Preserve auth/VPN logs (success+fail) + NAT mappings for the chained window: pivot-ready [Evidence source: FortiGate] ▪ Snapshot exposure state ("what was internet-facing, when, why"): anchor CTEM remediation [Evidence source: FortiGate / FortiCNAPP / FortiRecon]
Artifacts to Extract	▪ Source IP/ASN clusters, scan fingerprints, usernames, ports/services: support clustering & scope [Evidence source: FortiGate; Enrichment: FortiRecon (optional)]
Reconstruction Guidance	▪ Rebuild timeline (scan/fingerprint → auth abuse) and document the "last safe point": handoff to remediation.

CISO: Risk · Exposure · Priorities · Decisions

Dimension	Dimension-Grade Insight
Primary Risk	Accelerated probability of initial access driven by industrialized exposure harvesting: threat recon→brute force as a leading risk indicator
Exposure Context	Internet-facing credential services and identity guardrails are systematically targeted: validate exposure ownership and posture continuously. [CTEM: FortiRecon]
CTEM Priorities	Continuous validation of external attack surface, credential services, and authentication controls: prioritize burn-down by criticality. [CTEM: FortiRecon + FortiCNAPP]
Strategic Decision	▪ Prioritize exposure reduction over reactive cleanup ▪ Fund identity hardening and automation ▪ Shift focus from alert volume to chained pre-intrusion signals

In this model, **weaponization** is not defined by creativity. It is defined by availability, packaging, and conversion velocity.

Weaponization: Industrialized Preparation and Adversary Enablement

Weaponization is no longer a “pre-attack” technical step. In 2025, it functioned as an industrial supply chain that converted raw inputs, credentials, access, exploits, tooling, and infrastructure into repeatable execution packages ready to activate once exposure is validated.

Fortinet’s visibility across dark web economy signals **(FortiRecon)** and email and browser execution telemetry **(FortiMail Workspace Security)** confirms the same operating model: **adversaries do not improvise**. They standardize, automate, and stockpile capabilities to reduce friction during exploitation and compress time-to-impact.

> Weaponization now begins with identity. When credentials and session material are available in bulk, exploitation becomes less dependent on vulnerability breakthroughs and more dependent on how quickly access can be operationalized."

Credential supply chains: stealer logs as industrial inventory (FortiRecon intelligence)
FortiRecon intelligence shows that credential theft is not a byproduct of attacks. It is an upstream industry producing inventory at scale. In 2025, FortiRecon recorded **4.62 billion stealer logs**, underscoring that credential harvesting now operates as a high-throughput production layer feeding multiple downstream economies, including initial access brokerage, phishing and impersonation, session hijacking, fraud, and ransomware staging.

Within dark web "database" activity, stealer logs dominated advertised and shared datasets **(67.12%)**, far exceeding other categories such as **combolists (16.47%)** and **leaked credentials (5.96%)**.

This distribution matters operationally. Stealer logs reduce attacker effort by bundling identity material with contextual artifacts, often including browser-resident data, enabling immediate replay and faster conversion than brute-force or password spraying.

MITRE ATT&CK alignment: Resource Development (TA0042) and identity enablement, including Acquire Infrastructure/Capabilities (T1583/ T1588) and credential-driven access pathways.

Inference: The prevalence of stealer logs strongly supports downstream **Valid Accounts (T1078)** and token or cookie-based access patterns, although specific token or cookie indicator values are not provided in the dataset.

Phishing and browser execution as industrial delivery: linkless, QR, and token-theft workflows with FortiMail Workspace intelligence

FortiGuard intelligence shows that modern weaponization is no longer limited to payload delivery. It increasingly focuses on execution pathways that minimize user friction and bypass traditional email link and attachment controls.

These behaviors reflect an industrial design requirement. Attackers optimize for repeatable execution under inspection. When a technique is reliably blocked by a control, the control changes the wrapper while keeping the underlying workflow intact. This mirrors the same logic observed in exploitation: vectors are interchangeable, while execution patterns are standardized.

IoC classes/pivots (no specific values provided):

- Email artifacts associated with QR delivery (embedded QR payloads, QR destination transitions)
- Indicators of multi-hop delivery (redirect chains, short-lived intermediate infrastructure)
- OAuth abuse indicators (consent prompts from anomalous contexts, unusual authorization flows)
- HTML smuggling indicators (HTML artifacts used for client-side reconstruction of payload delivery)

ATT&CK mapping (partial, based on described behaviors):

- **Phishing (T1566)** as the industrial delivery surface
- **Steal Application Access Token (T1528)** as the identity acceleration pattern (OAuth token theft)
- **Drive-by compromise / browser-based execution** patterns are plausible but not confirmed beyond the described techniques in the intake

Across patterns captured in Perception Point intake, weaponized email operations emphasized:

- **Linkless payload delivery**, reducing dependency on URL reputation controls
- **Multi-stage redirect and delivery chains**, using operationally interchangeable infrastructure layers
- **QR-code–based delivery flows**, shifting execution from email scanners to mobile and browser contexts
- **OAuth token theft and account takeover workflows**, weaponizing trust and identity rather than malware
- **HTML smuggling** as an execution-enablement mechanism, shifting payload assembly into the client context

Vulnerability Commoditization: Exploits as Stock, Not Events

Defensive velocity is the only control that scales.

FortiRecon dark web intelligence shows that vulnerability discussion and exploitation enablement in 2025 was not primarily about "new CVEs." It was about sustained market attention on vulnerabilities that are operationally useful, widely deployed, reliably exploitable, and compatible with automation. In the dark web vulnerability landscape, FortiRecon recorded **656 CVEs discussed**, indicating that vulnerability-driven capability sourcing is continuous rather than episodic.

The weaponization takeaway is that CVEs become "industrial" when they are sufficiently packaged with scripts, modules, guides, proof code, and operational playbooks, so exploitation can run as a repeatable loop rather than a bespoke intrusion. That packaging is what converts disclosure into execution at scale.

This is where "vulnerability readiness" changes the defender's definition of exposure. Darknet vulnerability discussions shift exposure from theoretical risk to a time-bound condition determined by exploit availability.

In 2025, FortiRecon observed **656 vulnerabilities** actively discussed on the darknet. Within that set, **344 (52.44%)** had publicly available PoC exploit code, **176 (26.83%)** had working exploit code, and **149 (22.71%)** had both PoC and working exploit code available. FortiRecon also observed **224 (34.15%)** already exploited in publicly reported threat campaigns. Once a vulnerability enters this ecosystem, especially when it includes working code and campaign evidence, exposure is no longer constrained by the attacker's capabilities. It is constrained by how quickly exploit material can be operationalized across targets.

Dark web vulnerability readiness

Metric	Value	What it indicates
Vulnerabilities actively discussed	656	"Exposure set" being operationalized in real time
With public PoC code	344 (52.44%)	Lower barrier to exploitation; faster packaging
With working exploit code	176 (26.83%)	Execution-ready material; repeatable exploitation loops
With both PoC + working code	149 (22.71%)	Highest readiness for scaled operationalization
Exploited in reported campaigns	224 (34.15%)	Evidence of real-world conversion and reuse

Notable zero days advertised in darknet discussions disproportionately affect widely deployed platforms and ecosystems, including Windows, Android, iOS, the GitHub API, AnyDesk, Chromium, and Cloudflare, reinforcing an industrial attacker's priority of maximizing footprint and reuse.

Defensive velocity is the only control that scales. When exploits become stock, the organizations that survive are those that can patch, detect, and contain **before weaponized exploitation spreads across their environment**.

AI as a force multiplier: tooling that compresses skill and time with FortiRecon intelligence

FortiRecon dark web signals also captured **AI-enabled offensive tooling** advertised as services and products, indicating a continued shift toward **capability abstraction** that reduces operator skill requirements and increases workflow speed.

A growing set of AI-enabled offensive tools is advertised on darknet forums and messaging channels, including **WormGPT (Official), FraudGPT, HexStrike AI, APEX AI,** and **BruteForceAI.**

These tools are explicitly marketed to reduce the effort required of attackers and accelerate reconnaissance, credential abuse, brute-force attacks, and attack-path generation. While they do not create new exposure, they compress the time required to validate and activate existing exposure, widening the velocity gap between attackers and defenders.

Automation is no longer limited to execution. It now governs exposure harvesting itself.

> “Automation is no longer limited to execution. It now governs exposure harvesting itself.”

FraudGPT and WormGPT	These AI-powered text generators help cybercriminals craft compelling phishing emails, fake business communications, and fraudulent legal documents. Unlike ChatGPT, these tools have no ethical restrictions, allowing attackers to refine scams, generate malicious code, and conduct social engineering at scale.
HexStrike AI	HexStrike AI is an underground offensive AI tool advertised to assist with automated reconnaissance, attack-path generation, and malicious content creation. It is positioned as an uncensored, attacker-oriented AI assistant designed to accelerate end-to-end offensive operations, from initial access to post-exploitation.
APEX AI	APEX AI is a self-hosted offensive security AI tool marketed for APT-style attack simulation, combining automated OSINT, attack chaining, and kill-chain generation to model end-to-end compromise paths up to ransomware deployment. It supports autonomous APT simulation and pentest modes aligned with OWASP/PTES, emphasizing uncensored analysis and full attacker tradecraft emulation.
BruteForceAI	BruteForceAI is an advanced penetration testing tool that integrates large language models (LLMs) for intelligent form analysis. It automatically identifies login form selectors using AI, then executes sophisticated multi-threaded attacks with human-like behavior patterns.

FortiGuard SecOps action box: weaponization

In an industrialized threat landscape, weaponization is not a preparatory step; it is a production system. Attacks are no longer improvised. They are assembled, packaged, and staged long before exploitation occurs.

SOC: Signals · Correlations · Detections · Automation

Phase	SecOps-Ready Actions
Signals to Monitor	▪ Stealer-log exposure referencing corporate domains (credentials + browser/session artifacts): treat as active exposure, not "intel." [Intel: FortiRecon] ▪ Access-for-sale listings (RDP/VPN/admin panels/server access): map listing type to your exposed surface and owners. [Intel: FortiRecon + Asset Context: FortiCNAPP] ▪ QR-based/linkless payloads + HTML smuggling indicators: flag weaponization built to bypass URL reputation. [Detection: FortiMail Workspace]
Correlation Logic	▪ Stealer-log exposure → valid authentication attempts in a compressed window: link exposed identity to login telemetry. ▪ Email execution signal → OAuth token activity / takeover attempt: chain mail artifact to identity events. ▪ CVE chatter spike → scanning/probing against matching asset classes: match "tech in chatter" to "tech you run." [Correlation/Case execution: FortiSOC] [Optional intel enrichment: FortiRecon]
Automation & Response	▪ Enrich stealer logs → identity inventory → asset criticality: route to owner with priority. ▪ On confirmed exposure: rotate credentials + revoke sessions/tokens: contain trust, not malware. ▪ Tighten conditional access (geo/device/behavior) for exposed identities: shrink the trust window. ▪ Quarantine QR/HTML-smuggling campaigns at scale: remove weaponization channel. ▪ Elevate exposed identities/assets into CTEM burn-down: force remediation, not monitoring. [Orchestration/Case execution: FortiSOC]

DFIR: Evidence · Artifacts · Reconstruction (compact)

Area	DFIR-Ready Guidance
Priority Evidence	▪ Stealer-log datasets tied to victim domain: preserve exposure proof + timestamps. ▪ Access-for-sale ads matching victim stack: preserve listing context + claimed access type. ▪ Email artifacts (QR payloads/HTML smuggling/redirect chain): preserve original + analysis output. [Evidence/Intel source: FortiRecon, FortiMail Workspace, FortiSOC]
Evidence Chain (What to Pivot On)	▪ Credential harvest → validation → access packaging: tie identity exposure to auth telemetry. ▪ Email/browser execution → token enablement: tie campaign artifacts to identity events. ▪ CVE discussion → tooling availability → probing: tie "chatter" to scans on your asset class. [Case/timeline correlation: FortiSOC] [Optional intel enrichment: FortiRecon]
Reconstruction Objective	▪ Prove upstream preparation prior to exploitation and document repeatability (packaging + staging). [Case/timeline: FortiSOC]

CISO: Risk · Exposure · Priorities · Decisions

Dimensions	Decision-Grade Insight
Primary Risk	Weaponization succeeds when identities, access, and exploits are pre-packaged faster than defenders revoke trust.
Exposure Context	Credential reuse + external identity trust (SSO/OAuth) + widely deployed tech reduce attacker effort and time-to-impact.
CTEM Priorities	▪ Treat credential exposure as active risk, not passive intelligence. ▪ Prioritize assets aligned with access-for-sale types and stealer activity. ▪ Validate identity controls under real-world weaponization scenarios. [CTEM: FortiRecon + FortiCNAPP]
Strategic Decisions	▪ Fund identity-centric detection over malware-only controls. ▪ Invest in dark web exposure monitoring as preventive control. ▪ Measure time-to-revocation (sessions/tokens/creds) as an executive metric.

Exploitation: Intrusion at Scale—The Industrialization of Execution

IPS Intelligence, FortiEDR / MDR Intelligence, FortiRecon Intelligence

Where exposure and weaponization materialize into intrusion

In 2025, exploitation no longer behaved like a “phase” defined by attacker ingenuity. It behaved like an automated race condition. Once exposure is validated, execution runs as a time-bounded loop whose only question is whether the weakness has been remediated. FortiGuard telemetry showed a structural imbalance, where reconnaissance and brute-force activity became narrower and more selective, while exploitation accelerated, because cybercrime pipelines front-loaded discovery and weaponization and then executed continuously at machine speed, without human decision points.

This is why exploitation must be treated as a performance problem. IPS telemetry measures pressure at scale. EDR telemetry confirms that when that pressure converts into host-level execution, a CVE shifts from theoretical risk to operational fact. In this model, exploitation is constrained less by the attackerss capability and more by the defenders velocity. The window between disclosure, exposure validation, and remediation defines success or failure. When MTTR exceeds attacker automation speed, intrusion is not just a possibility. It is the expected outcome.

Globally, exploitation pressure in 2025 was not campaign-driven. It followed an industrial cadence (Table A). The signal is consistent. Pressure persists until exposure is removed or execution succeeds. Regionally, conversion speed and growth diverge (Table A), reflecting differences in exposure density, patch cadence, and technology mix. North America showed the highest absolute volume and fastest acceleration, while LATAM showed lower absolute volume but faster growth, consistent with rising exposure-to-exploitation conversion.

When MTTR exceeds attacker automation speed, intrusion is not just a possibility. **It is the expected outcome.**

> When MTTR exceeds attacker automation speed, intrusion is not just a possibility. It is the expected outcome."

Table A: At-a-Glance (Volumes / Velocity / Distributions)

Metric	2025 value	Notes
Global exploitation attempts (IPS)	~121.99B	+25.49% YoY
Sustained cadence	~334M/day; ~13.9M/hour; ~3,850/sec	Continuous
North America	~54.29B	YoY +173.47%
APAC	~32.44B	YoY −21.05%
EMEA	~20.36B	YoY −20.51%
LATAM	~14.91B	YoY +39.85%

Metric	2025 value	Notes
Newly exploited vulns by release year	2025: 212 (58.24%); 2024: 74 (20.33%); 2023: 20 (5.49%); 2022: 13 (3.57%); 2021: 7 (1.92%)	Novel attacks is the new normal
Exploited vulns by actor mix	APT: 124 (19.53%); Ransomware: 145 (22.83%); Both: 130 (20.47%); Unknown: 274 (43.15%)	Symbiotic relationship between APT and cybercrime grows
Exploit availability	PoC: 342 (53.86%); Weaponized: 198 (31.18%); Both: 158 (24.88%)	Entering the industrialized exploit age
SharePoint execution spike (MDR)	4,682 confirmed malicious executions	Week of July 07; CVE-2025-49706 / CVE-2025-49704
EDR signal concentration	Suspicious Application 48.96%; Malicious File Detected 12.91%; Sensitive Information Access 11.50%; Suspicious Script Execution 4.42%; Process Hollowing/Injection ~8% (combined)	Execution + Defense Evasion dominant
Fileless PowerShell signal	Generic.powershell.fileless: 10.46%	Post-exploitation execution signal

What drives impact is no longer how old a vulnerability is, but how quickly it becomes automation-ready (Table A). Among newly exploited vulnerabilities observed in 2025, **212 (58.24%)** originated from 2025 releases, followed by **74 (20.33%**) from 2024, **20 (5.49%)** from 2023, **13 (3.57%**) from 2022, and **7 (1.92%**) from 2021. This velocity bias is amplified by exploit availability: **342 (53.86%)** have public PoC exploit code, **198 (31.18%)** have public working weaponized exploit code, and **158 (24.88%**) have both (Table A). Once weaponized code becomes public, exploitation shifts from possibility to probability, because success is constrained less by skill and more by access to automation-ready material.

The actor ecosystem reinforces the same pattern—exploitation throughput is broadly distributed and increasingly commodity-scale (Table A). Publicly reported campaigns tie **124 (19.53%)** vulnerabilities to APT groups, **145 (22.83%)** to ransomware groups, **130 (20.47%**) to both APT and ransomware groups, and **274 (43.15%)** to unknown or unattributed actors. In an industrialized ecosystem, attribution matters less than throughput. Shared infrastructure, shared tooling, and shared exploit availability allow impact to be driven at scale, often without a named actor.

When exploitation succeeds, EDR telemetry shows it primarily manifests as **Execution** and **Defense Evasion**, not as novel malware delivery (Tables A/B). Dominant 2025 triage signals concentrate on anomalous use of legitimate executables and native mechanisms: **Suspicious Application (48.96%) (ATT&CK: T1059; T1218), Malicious File Detected (12.91%) (ATT&CK: T1204), Sensitive Information Access (11.50%) (ATT&CK: TA0006; TA0007), Suspicious Script Execution (4.42%) (ATT&CK: T1059.001; T1059.003**), and **Process Hollowing/Injection (~8% combined**) (ATT&CK: **T1055**). This reframes any investigation. Not "Was malware dropped?" but "Did unauthorized execution occur in a way designed to evade controls?" In this model, **Execution (TA0002)** and **Defense Evasion (TA0005)** are not post-compromise refinements; they are design requirements of exploitation.

A concrete illustration appears in the early July surge linked to the large-scale exploitation of Microsoft SharePoint vulnerabilities **CVE-2025-49706** and **CVE-2025-49704**. During the week of July 7, MDR telemetry recorded **4,682** confirmed malicious executions, an order-of-magnitude increase over adjacent periods (Table A). The transition did not ramp gradually. Once exposure aligned with weaponized capability, execution was activated immediately and at scale. Operationally, the exploit vector was the trigger, but the execution logic was consistent: **.NET binaries are loaded directly into memory**, no traditional malware files required. The same pattern recurred throughout the year via multiple upstream conditions: alternative web application vulnerabilities, compromised machine keys, ViewState abuse, and webshell-assisted execution (Table B).

The takeaway is structural. The exploit vector is interchangeable; the execution layer is standardized to maximize reuse, automation, and speed across heterogeneous environments.

This is also visible in what attackers choose to exploit. Telemetry shows concentration, not diversification: **web application remote code execution** (including Log4j-class behaviors, SharePoint flaws, and framework deserialization issues), **SMB exploitation** used as a bridge into lateral movement, **HTTP authentication abuse** against exposed administrative interfaces, and **IoT/edge vulnerabilities** (notably routers and cameras) where persistent exposure and weak patch hygiene preserve reach (Table B). Selected vulnerabilities consistently meet three operational requirements: a broad deployment footprint; stable, well-tested exploit code; and seamless compatibility with automation and LOLbins. Post-exploitation execution signals reinforce the same principle, including **PowerShell-based fileless execution (generic.powershell.fileless: 10.46%)** and extensive reliance on legitimate remote administration and tunneling utilities, such as **AnyDesk, LogMeIn, Radmin, Atera, Ngrok, WinExe, and WinVNC variants** (Table B).

Key insight: Attackers are no longer optimizing for how advanced their malware is. They are optimizing for how consistently and how fast execution succeeds.

"Old" vulnerabilities persist not because attackers prefer the past, but because automation favors what is predictable and repeatable. Exploitation persistence correlates more strongly with **patch latency** and **exposure duration** than with novelty. Once a vulnerability is integrated into automated pipelines, it does not age out as long as exposure remains. FortiGuard telemetry aligns persistence with vulnerabilities already operationalized in attacker frameworks: those with confirmed exploitation in the wild (as reflected in **CISA Known Exploited Vulnerabilities)**, sustained exploit probability (high **EPSS** scores), and public exploit availability and tooling reuse. This explains why vulnerabilities such as **Log4j**, **SMB-related flaws**, and **IoT management CVEs** continue to dominate years after disclosure. Proven vulnerabilities introduce velocity; novelty introduces uncertainty. Industrialized exploitation is a stopwatch problem. If exposure is not removed and execution is not interrupted faster than automation can convert it, intrusion stops being a risk and becomes the default.

Table B: Targets / TTPs / Services / Tools

Layer	What was exploited / observed	Evidence signals / examples
Exploitation concentration	Web app RCE (Log4j-class behaviors, SharePoint flaws, deserialization); SMB exploitation; HTTP authentication abuse; IoT/edge vulns (routers, cameras)	Standardization beats innovation selection for reach + reliability + automation compatibility
Execution + Defense Evasion (EDR)	T1059; T1218; T1204; T1059.001; T1059.003; T1055; plus TA0006 and TA0007 indicators	Signal mix: Suspicious Application, Script Execution, Process Injection, Sensitive Info Access
SharePoint execution pattern (MDR)	CVE-2025-49706 / CVE-2025-49704 triggering a standardized execution layer	.NET binaries loaded in memory; no traditional malware files; recurring triggers: alternative web app vulns, compromised machine keys, ViewState abuse, webshell-assisted execution
Living-off-the-land & admin tooling	AnyDesk, LogMeIn, Radmin, Atera; Ngrok; WinExe and WinVNC variants	Abuse of legitimate tooling and native execution mechanisms
"Old vuln" persistence drivers	KEV status; high EPSS; public exploit availability + tooling reuse	Examples referenced: Log4j; SMB-related flaws; IoT management CVEs

Time-to-exploit: from disclosure to exploitation at industrial speed

FortiGuard Labs analyzed 2025 CVEs using **Critical Outbreak Alerts** to measure TTE, defined as the elapsed time between public CVE disclosure and the first exploitation attempt observed in global Fortinet telemetry. TTE values are day-normalized to account for global telemetry timing and update cycles.

The results show a clear and consistent compression of TTE. Across the cases reviewed, the window between vulnerability disclosure and first observed exploitation ranged from zero to only a few days. In numerous critical outbreaks, exploitation activity appeared on the same or the following day. This indicates that attackers are no longer treating new vulnerabilities as items to be evaluated over time, but as inputs to automated workflows that can be operationalized immediately once exploit material becomes available.

This behavior is not simply opportunistic. It reflects the industrialization of cybercrime, where vulnerabilities are rapidly integrated into automated exploitation pipelines and executed at scale as soon as they become usable.

Executive Table: TTE (day-normalized)

Outbreak Alert	CVE(s)	CVE Date	TTE (days)
Fortra GoAnywhere MFT Attack	CVE-2025-10035	18/09/25	0
Oracle E-Business Suite RCE Zero-day	CVE-2025-61882	05/10/25	0
React2Shell Remote Code Execution	CVE-2025-55182; CVE-2025-66478	03/12/25	1
Cisco ASA and FTD Firewall RCE	CVE-2025-20333; CVE-2025-20362; CVE-2025-20363	25/09/25	1
Apache Tomcat RCE	CVE-2025-24813	10/03/25	0

FortiGuard SecOps action box: exploitation

In an industrialized threat landscape, exploitation is not stopped by awareness; it is stopped by speed. Detection, patching, and response must operate faster than automated execution pipelines.

SOC: Signals · Correlations · Detections · Automation

Exploitation Phase

Phase	SecOps-Ready Actions
Signals to Monitor	▪ Repeated exploitation attempts against public-facing services (web apps, SMB, edge devices). [Detection: FortiGate IPS] ▪ IPS detections tied to KEV-listed CVEs and high-EPSS vulnerabilities [Detection: FortiGate IPS] ▪ Web services spawning interpreters (PowerShell, cmd, .NET loaders) [Detection: FortiEDR] ▪ Fileless execution indicators and LOLbin usage following network alerts [Detection: FortiEDR]
Correlation Logic	▪ IPS exploitation alert → EDR execution on the same asset within short time window ▪ Repeated exploitation attempts against assets with known patch latency ▪ Execution behavior recurring across multiple hosts after similar exploit signals [Correlation/Case execution: FortiSOC]
Detection Gating (Operational)	▪ Low: Single exploit attempt without execution ▪ Medium: Repeated exploitation attempts without confirmed execution ▪ High: Exploitation + host-level execution (fileless, script-based, or injected)
Automation & Response	▪ SOAR correlation: IPS exploit → EDR execution → asset criticality ▪ Automated host isolation on High severity execution signals ▪ Implement from trigger and exposure remediation workflows [Orchestration/Case execution: FortiSOC]

DFIR: Evidence · Reconstruction · Proof of Execution

Exploitation → Execution Transition

Area	DFIR-Ready Guidance
Priority Evidence	▪ IPS exploitation logs tied to specific CVEs or signatures ▪ EDR execution telemetry (fileless scripts, LOLbins, injected processes) ▪ Process trees showing execution from web services or network-facing daemons [Evidence source: FortiEDR, FortiGate] [Case/timeline correlation: FortiSOC]
Evidence Chain (What to Pivot On)	▪ Exploit attempt → process creation → in-memory execution ▪ Timing alignment between IPS alerts and EDR execution events ▪ Reuse of identical execution techniques across different assets
Expected Behaviors	▪ No malware dropped to disk ▪ PowerShell or .NET execution immediately after exploitation ▪ Legitimate binaries used for execution and persistence
Reconstruction Objective	▪ Prove transition from exploitation to execution ▪ Demonstrate automation and repeatability of execution ▪ Attribute intrusion to industrialized exploitation rather than opportunistic attack

CISO: Risk · Exposure · Priorities · Decisions

Exploitation as a Speed Problem

Dimension	Decision-Grade Insight
Primary Risk	Exploitation succeeds when attacker automation outpaces patching and response velocity
Exposure Context	Internet-facing services with patch latency and inconsistent hardening are repeatedly exploited
CTEM Priorities	▪ Continuous validation of exposed services ▪ Prioritize KEV + high-EPSS vulnerabilities ▪ Measure and reduce patch latency as a security control [CTEM: FortiRecon]
Strategic Decisions	▪ Fund automation over manual remediation ▪ Treat patch latency as business risk ▪ Shift KPIs from vulnerability counts to execution prevention


Command-and-control is not an event.
It is a background process at scale.

The signal is structural. Cybercrime is not operating through isolated campaigns, but as a **persistent service model** in which millions of devices maintain continuous contact with criminal infrastructure.

Post-Exploitation: When Cybercrime Takes Control at Machine Speed

This is the phase where technical compromise becomes data loss, business disruption, and regulatory exposure, and where failure directly determines breach cost and recovery time.

In 2025, post-exploitation no longer functioned as a reactive follow-up. It operated as the industrial core of modern cybercrime, a standardized internal execution layer engineered to convert access into control through automation, speed, and repeatability.

Network detection and response (NDR) intelligence shows that once a foothold exists, adversaries do not pause to reassess. They activate a pre-built pipeline where persistent command-and-control is established, stolen credentials are operationalized, discovery is automated, and lateral movement propagates through trusted administrative paths. Domain controllers are pressured early, and when conditions align, operations shift seamlessly into data theft or ransomware staging, often before defenders can move from initial detection to coordinated response. In this model, post-exploitation is constrained primarily by defensive velocity. If correlation, containment, and remediation cannot keep pace with automation, lateral expansion and domain compromise become expected outcomes rather than edge cases.

Control at scale begins with stealthy connectivity built for continuity. Adversaries increasingly favor TCP-based beaconing across standard and non-standard ports, while continuing to leverage SSL and HTTP to blend into enterprise flows. Beaconing is no longer just communication. It functions as the control layer for automated operations, running discovery, lateral movement, and remote execution in parallel. NDR observed beaconing becomes high-confidence post-exploitation when correlated with newly seen NTLM authentication deviations, remote

Table C: Signals / Services / Tools / Protocols

Layer	What was observed	Signals / examples
C2 / control layer (NDR)	TCP-based beaconing on standard & non-standard ports; SSL & HTTP blending	High confidence when correlated with newly seen NTLM deviations, remote command execution, early probing of domain controllers
Botnet infrastructure (FortiGate)	C2 communications with known botnet infrastructure	Confirms compromise + active participation in criminal operations
Botnet families (regional examples)	APAC: Mirai, Mozi, Prometei, Morto	Optimized for large-scale infrastructure abuse across IoT, edge devices, exposed services
Botnet families (LATAM)	Phorpiex, Prometei	Sustained presence rather than short-lived spikes
Botnet families (EMEA)	Zeus, Mozi, Andromeda, Prometei	Reduced noise, preserved capability
Botnet families (North America)	SystemBC, Mirai, Morto	Associated with C2 relay, initial access brokerage, post-compromise coordination
Identity-driven movement (NDR)	Abuse of legitimate admin tooling for lateral movement & remote execution	Impacket, WinRM, WMIC, scheduled task mechanisms; “newly seen” NTLM patterns
Automated discovery (NDR)	High-volume RPC + named pipe abuse	wkssvc used to enumerate users/devices
Accelerators (NDR)	Technical debt + trusted tool blind spots	Legacy protocol SMBv1; misuse of RMM tools
Impact indicators (NDR)	Shift from expansion to theft/ransomware staging	Large outbound transfers, multi-threaded exfiltration, SSH/FTP to rare/cloud/new destinations

command execution, and early probing of domain controllers (Table C). The intent is operational. Attackers do not need perfection, only enough speed to outrun defenders.

FortiGate C2 intelligence provides a direct view of this control layer at an industrial scale. These detections are signature-based triggers that occur when infected devices establish command-and-control communications with known botnet infrastructure, confirming both compromise and active participation in criminal operations.

Across 2025, FortiGate sensors recorded **7.10 billion botnet detections** globally, equivalent to approximately **592 million per month** and **19.4 million per day** (Table D). The signal is structural. Cybercrime is not operating through isolated campaigns, but as a persistent service model in which millions of devices maintain continuous contact with criminal infrastructure. Command-and-control is not an event. It is a background process at scale.

> The cost of a breach is set after entry by how quickly we stop lateral spread and data movement."

Telemetry shows how automation manifested differently by region: **APAC** exhibited the highest scaling velocity, **LATAM** showed sustained growth with country-level concentrations, **EMEA** showed reduced noise with preserved capability, and **North America** showed lower, but more selective use associated with control and monetization (Table D). Across all regions, long-lived botnets, including **Phorpiex, Prometei, Mirai, Mozi**, and **SystemBC**, functioned as an industrial infrastructure that sustains external control, enables lateral operations, relays traffic, and supports monetization workflows.

Inside the environment, speed comes from trust. NDR consistently observed the abuse of legitimate administrative tools, **Impacket**, **WinRM**, **WMIC**, and **scheduled tasks**, to automate lateral movement and remote command execution (Table D). Operations increasingly relied on credential theft rather than brute force, with **"newly seen" NTLM patterns** and anomalous authentication activity emerging as primary network-level indicators of compromise. In this model, identity becomes a distributed acceleration layer: valid credentials allow attackers to traverse environments, execute commands, and access sensitive services while remaining operationally legitimate. The defender challenge is less about visibility than timing. Many organizations lack the correlation maturity to recognize legitimate trust being operationalized faster than they can respond.

Reconnaissance also runs at machine speed. NDR observed high-volume RPC activity and the systematic abuse of named pipes, specifically **wkssvc**, to enumerate users and devices in automated workflows. In representative cases, a single account was associated with more than **100 internal IP addresses** and repeatedly accessed **wkssvc** (Table D).

Table D: At-a-Glance (Volumes / Regional Intensity / Observed Thresholds)

Metric	2025 value	Notes
Global botnet detections (FortiGate sensors)	7.10B	Botnet detections confirm active C2
Global C2 cadence	~592M/month; ~19.4M/day	Every day of the year
APAC botnet detections	3.25B	+26.21% YoY; ~270M/month; ~8.9M/day
LATAM botnet detections	2.09B	+12.05% YoY; ~174M/month; ~5.7M/day
EMEA botnet detections	1.22B	−57.60% YoY; ~102M/month; ~3.3M/day
North America botnet detections	538.31M	−27.09% YoY; ~44.9M/month; ~1.47M/day
LATAM concentration (countries)	Mexico; Panama; Venezuela; Brazil; Colombia	Heavily concentrated in Mexico, followed by listed countries
Representative discovery threshold	>100 internal IPs	Single account repeatedly accessing wkssvc

This is machine-driven discovery. Once enumeration completes, attackers rapidly identify domain controllers, map privilege relationships, and prioritize high-value targets. Across investigations, detection at enumeration was often the last meaningful time advantage, when response windows were measured in hours rather than minutes, before workflows progressed toward irreversible impact.

Expansion turns into impact when enterprise accelerators align. NDR identified two structural amplifiers: **technical debt** and **blind spots in trusted tools**. The continued presence of legacy protocols such as **SMBv1** creates predictable, low-effort lateral movement paths that integrate seamlessly into automated workflows.

In parallel, the growing misuse of **Remote Monitoring and Management (RMM)** tools enables persistent, high-privilege access that blends into routine IT activity unless strong baselining and cross-signal correlation are in place (Table B).

When these conditions align, detections associated with confirmed ransomware operations or data theft include **large outbound data transfers**, **multi-threaded exfiltration**, and **SSH** or **FTP** activity to rare, cloud-based, or newly observed destinations, typically alongside command-and-control, lateral movement, and credential abuse (Table B). In several investigations, exfiltration began while discovery and lateral movement were still active, further compressing response windows.

By the time these network-level indicators surface, the internal operation has reached operational maturity. This is why NDR emphasizes early warning detection, **authentication anomalies**, **automated enumeration**, **beaconing, vulnerable protocol usage**, **and suspicious RMM activity**, well before malicious executables are observed on endpoints.

These patterns were observed predominantly in complex enterprise environments and may vary across smaller or less mature infrastructures.

The cost of a breach is set after entry by how quickly we stop lateral spread and data movement.

FortiGuard SecOps action box: post-exploitation

In an industrialized threat landscape, post-exploitation is not stopped by visibility. It is stopped by velocity. Control, expansion, and impact unfold as automated workflows. Defensive success is defined by how quickly SecOps can detect, correlate, and contain them.

SOC: Signals · Correlations · Detections · Automation - Post-Exploitation (NDR + Botnet C2)

How to read this table:

- ***Baseline actions =*** *achievable by mid-maturity SOCs to get visibility + triage*
- ***Machine-speed actions =*** *correlation + automated containment*

Phase	SecOps-Ready Actions
Signals to Monitor	Baseline: ▪ Confirmed C2 callbacks (Botnet intelligence hits): treat as compromise signal, not "suspicious." ▪ Persistent outbound cadence (beaconing): baseline periodicity per host/destination. ▪ Authentication anomalies (IAM/AD): new logon types, unusual hosts/users. [Detection: FortiGate] Machine-speed: ▪ Beaconing over TCP/SSL/HTTP across common/uncommon ports: detect protocol/port drift + periodicity. ▪ Newly seen NTLM/Kerberos patterns vs baseline: treat as expansion indicator. ▪ Automated discovery signals (wkssvc / high east–west fan-out): identify burst discovery + lateral prep. ▪ Rare protocols/destinations (SSH/FTP to cloud/new domains): flag new egress paths used for staging/movement. [Detection: FortiNDR]
Correlation Logic	Baseline: ▪ C2 + same-host auth anomaly: C2 + identity signal = active control. ▪ Repeated beaconing + lateral protocol usage: beacon + SMB/RPC/WinRM adjacency. Machine-speed: ▪ C2 → identity abuse → discovery → lateral movement overlapping (not linear): correlate concurrent signals into one incident. ▪ Compressed timelines between phases: escalate when transitions happen in minutes. [Correlation/Case execution: FortiSOC]
Automation & Response	Baseline: ▪ Trigger containment playbook on confirmed C2: contain first, investigate after. ▪ Rapid credential reset + privilege review: stop credential reuse. ▪ Segment/block at gateway: break C2 and cut expansion paths. Machine-speed: ▪ Correlate NDR + Botnet + IAM + EDR as one case: one queue, one owner. ▪ Auto-isolate hosts on confirmed C2 callbacks: containment without ticket latency. ▪ Automated credential containment (sessions/privileges): reduce identity-abuse dwell time. ▪ Auto-classify as "Machine-speed Post-Exploitation": immediate escalation posture. [Orchestration/Case execution: FortiSOC]

DFIR: Evidence · Reconstruction · Proof of Machine-speed Execution

Area	DFIR-Ready Guidance
Priority Evidence	▪ Botnet C2 hits + callback cadence (prove control). ▪ Network flow metadata for beaconing + east–west fan-out (prove automation). ▪ Authentication logs showing new/rare identity usage (prove abuse). ▪ SMB/RPC/WinRM traces indicating lateral movement (prove expansion). [Evidence source: FortiNDR] [CTEM: FortiRecon]
Evidence Chain	▪ Baseline: C2 → suspicious auth → lateral movement (minimum proof chain). ▪ Machine-speed: C2 → identity abuse → automated discovery → expansion → data movement (compressed, overlapping phases). [Case/timeline correlation: FortiSOC]
Reconstruction Objective	▪ Prove compromise was active (not attempted) and show automation + compressed timelines as the impact driver.

CISO: Post-Exploitation as a Business Velocity Risk

Area	Decision-Grade Insight
Primary Risk	Attackers move faster internally than teams can detect and coordinate: post-exploitation is a business velocity risk.
Exposure Context	Identity sprawl, legacy protocols, and trusted admin tooling accelerate attacker speed and reach.
CTEM Priorities	Baseline: validate botnet C2 as breach confirmation; measure detection-to-containment time. Machine-speed: validate post-exploitation paths end-to-end; reduce identity abuse dwell time. [Detection: FortiGate, FortiNDR, FortiEDR; Execution: FortiSOC]
Strategic Decisions	▪ Invest in correlation + response speed before adding tools. ▪ Treat botnet C2 as breach confirmation. ▪ Shift KPIs from alert volume to time-to-containment.

In an industrial threat model, single anomalies are noise. **Sequences are intent.**

Industrialized Cloud Intrusions: Identity, Automation, and Control at Machine Speed

FortiCNAPP Intelligence

Industrial cloud cybercrime does not exploit clouds. It weaponizes trust. Organizations must act at machine speed.

Cloud control plane abuse: when scale and velocity redefine intrusion

Cloud environments continue to experience a dual acceleration, where attack velocity and operational complexity are increasing simultaneously. As organizations scale cloud adoption, adversaries are shifting away from exploiting infrastructure vulnerabilities and toward abusing **identity and access pathways** and **misconfiguration gaps** to gain persistence, monetize resources, and commit fraud. **FortiCNAPP intelligence** confirms this shift.

This evolution reflects the industrialization of cloud cybercrime. Attacks are no longer driven by bespoke techniques or manual decision-making. They execute as repeatable operational workflows optimized for speed, low cost, and predictable outcomes.

FortiCNAPP addresses this reality by unifying posture management, workload protection, and identity threat detection into a **single cloud-native application protection platform (CNAPP)**. The strategic advantage is not broader visibility but **earlier confirmation of the attacker's intent**.

A defining capability underpinning this intelligence is the use of **composite alerts**, detections that correlate multiple weak signals (identity anomalies, API behavior, privilege changes, and infrastructure actions) into high-confidence indicators of active compromise. In an industrial threat model, single anomalies are noise. Sequences are intent.

Across 2025 telemetry, four patterns consistently defined cloud intrusions:

- **Identity compromise** remained the dominant intrusion vector.
- **Discovery-heavy activity** preceded most high-impact incidents.
- **Monetization** clustered around **SES abuse, resource hijacking,** and **cryptomining.**
- **Regional and industry context** materially affected both attack prevalence and detection confidence.

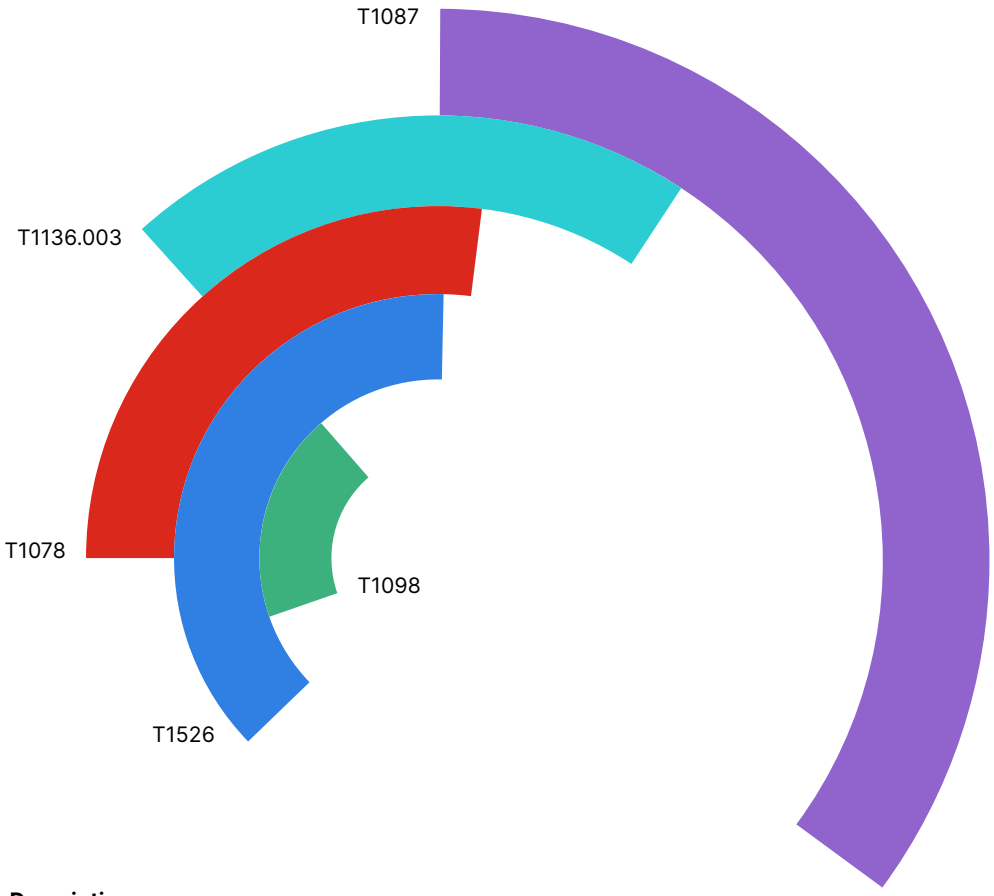

Technique	Description
T1078 – Valid Accounts	Core technique enabling cloud takeover
T1087 – Account Discovery	Enumeration of users and roles
T1526 – Cloud Service Discovery	Mapping cloud attack surface
T1098 – Account Manipulation	Persistence and privilege expansion
T1136.003 – Create Cloud Account	Long-term access establishment

Identity as the cloud control plane: industrialized access abuse

Identity is the control plane of the cloud. Once compromised, valid credentials allow adversaries to operate entirely within expected provider behavior, without malware, without exploits, and often without triggering traditional security controls. **FortiCNAPP intelligence** confirms that throughout 2025, the majority of confirmed cloud incidents originated from stolen, exposed, or misused credentials rather than from infrastructure exploitation. This enables living-off-the-land cloud attacks that blend into normal operations unless behavioral correlation is applied.

With valid credentials, attackers can bypass network-centric controls, abuse legitimate APIs for discovery and privilege expansion, blend malicious actions into operational baselines, and execute automation at a pace defenders struggle to match. This marks a fundamental shift. In cloud environments, **MITRE ATT&CK** no longer explains how attackers break in. It explains how they operate at scale. The control plane is no longer infrastructure. It is trust, and trust is enforced through identity.

Regional and sectoral context reinforced the same identity thesis. Identity-driven incidents were observed globally with meaningful regional differences. APAC showed the highest relative prevalence; AMER (US) showed high activity with stronger detection maturity; AMER, including LATAM, reflected mixed exposure and exploitation; and EMEA showed more balanced exposure aligned with detection capability.
An inverse relationship was observed between the prevalence of pen testing and successful identity compromise, indicating that proactive testing materially reduces attacker dwell time.

Host-based compromises remained consistent across regions, with 11–13% of environments experiencing at least one incident involving compromised hosts or workloads. This reinforces the conclusion that identity abuse, not endpoint exploitation, defines the dominant risk in the provided data.

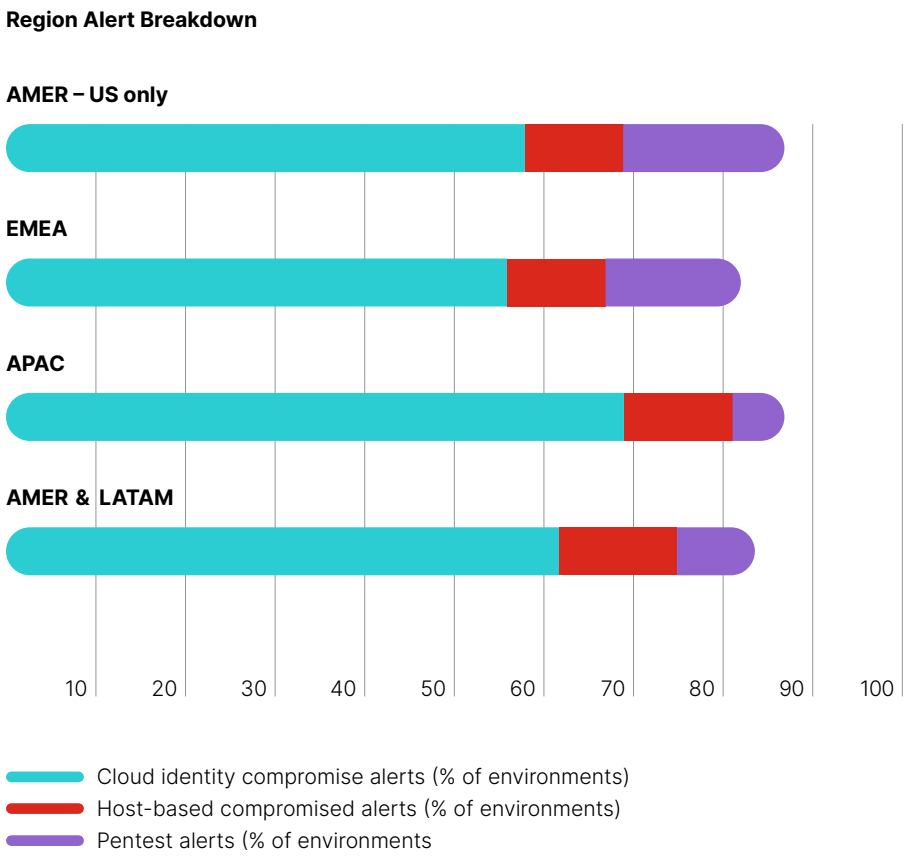

Figure 1: Compromises and pen testing by region

> The implication is that **identity sprawl, not infrastructure weakness, now defines cloud exposure.**”

Sector analysis reinforces the same conclusion. **M-adjacent organizations, retail, consulting**, and **communications** showed the highest concentration of identity compromise, driven by large identity populations, federated access models, and complex cloud integrations. **Technology, banking**, and **professional services** demonstrated moderate risk, while **education, hospitality**, and **smaller software segments** showed lower observed compromise rates, though this gap continues to narrow as cloud adoption accelerates.

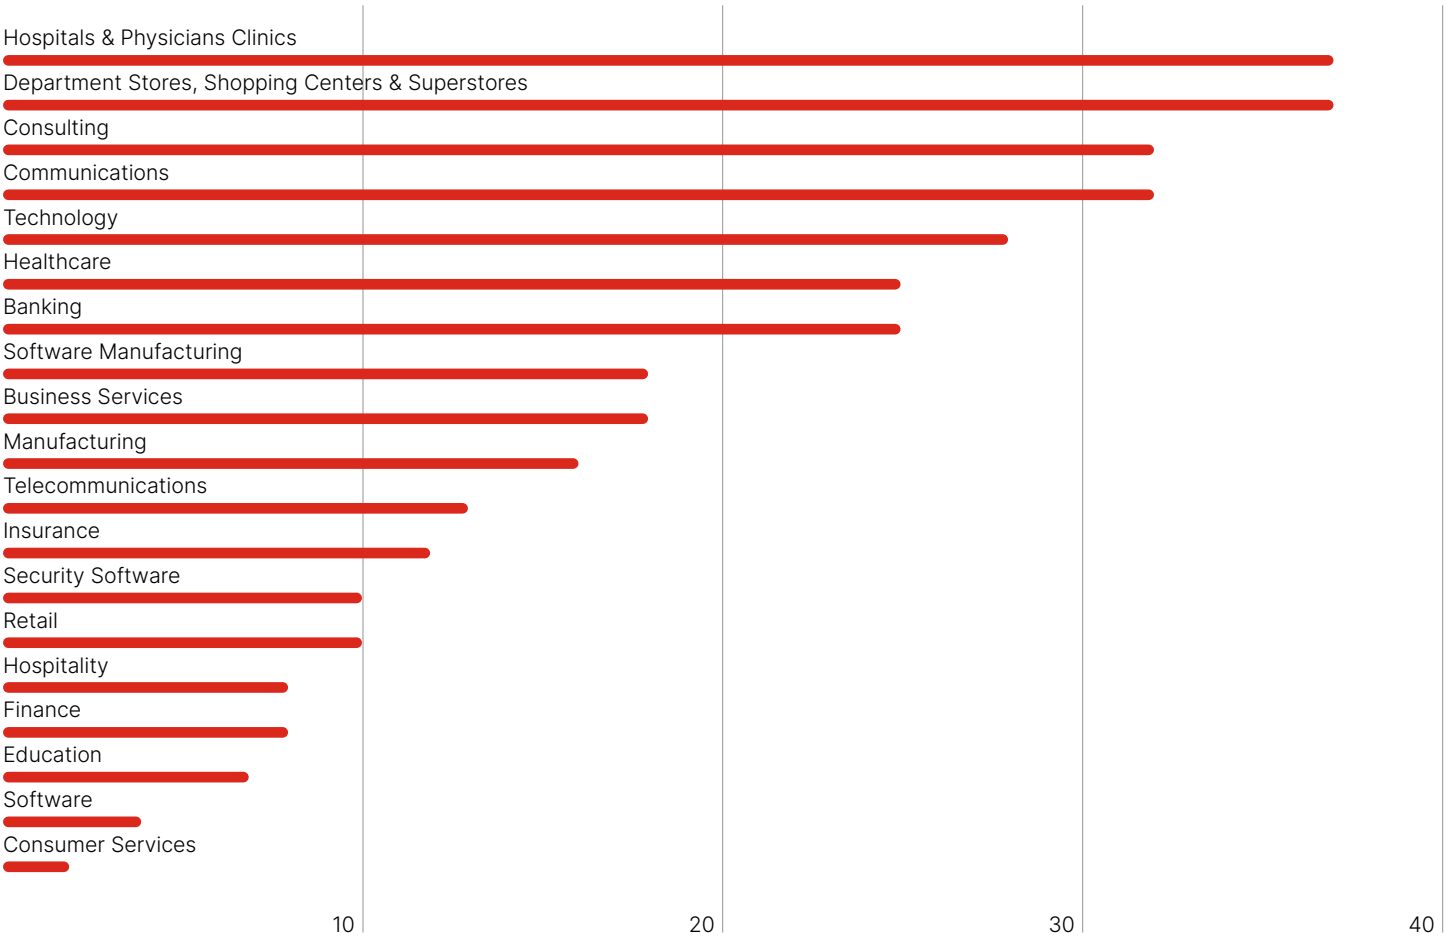

Figure 2: Identity compromise by sector

Discovery-first operations: how cloud cybercrime learns faster than defenders

A defining characteristic of cloud intrusions in 2025 was automated discovery at scale. **FortiCNAPP intelligence** shows that discovery-heavy API activity consistently preceded most high-impact incidents, acting as the pivot from initial access to operational control. In an industrialized cloud threat model, discovery is not a "next step." It is the workflow trigger that converts credentials into repeatable execution.

This discovery is not exploratory. It is machine-driven reconnaissance, executed in bursts across services, regions, and accounts to compress time-to-impact. Functionally, discovery becomes the attacker's industrial quality-assurance phase, rapidly profiling environments for reliable paths to scale:

- Cloud resources and services are enumerated.
- Identity relationships and permissions are mapped.
- Privilege escalation paths are validated.
- High-value targets are prioritized.

Once discovery automation is complete, investigation becomes retrospective rather than preventive. Containment windows collapse from hours to minutes, and defender advantage erodes as automated workflows dictate progression. Across DFIR investigations, this discovery phase consistently represents the last meaningful opportunity to interrupt an attacker's momentum before privilege escalation and monetization begin. **When the attacker moves at machine speed, slow response becomes a business risk**.

Exposure as an accelerator: misconfiguration, monetization, and impact

Misconfigurations remain one of the most persistent and consequential risk factors across cloud environments. While misconfigurations alone do not always constitute an incident, **FortiCNAPP intelligence** shows that they significantly amplify an attacker's velocity once identity is compromised. In industrialized cloud attacks:

- **Credentials are the trigger.**
- **Misconfigurations are the accelerator.**

Overly permissive IAM policies, long-lived access keys, unrestricted network access, publicly exposed storage, missing immutability, and root accounts without MFA reduce the effort required by attackers and expand the blast radius. These conditions do not simply increase risk. **They shorten the path from access to impact.**

Once discovery and privilege expansion are complete, adversaries transition into repeatable monetization workflows. Across 2025 telemetry, these consistently cluster around:

- **Cloud email service abuse (SES)**
- **Resource hijacking**
- **Cryptomining**

These actions are not improvised. They are execution stages of industrial cybercrime, designed to convert access into financial return with minimal delay. In this operating model, speed compensates for imperfection. Attackers do not require precision, only momentum.

The strategic implication is that any organization that treats cloud incidents as isolated security failures is misdiagnosing a structural problem in its operating model. Cloud incidents are no longer defined by how access is gained, but by how quickly trust is operationalized.

2025 FortiCNAPP actions: high severity

- ↘ **Rotate access keys every 350 days or less**
- ↘ **Ensure Network Access Control Lists (ACLs) do not allow unrestricted inbound traffic**
- ↘ **The security group attached to the Network Interface should not allow inbound traffic from all ports**
- ↘ **Use locked immutability policies for Storage Accounts Blobs with business-critical data**
- ↘ **The security group attached to an EC2 instance should not allow inbound traffic from all ports**
- ↘ **Ensure the attached S3 bucket policy does not grant "Allow" permission to everyone**
- ↘ **Enable multi-factor authentication (MFA) for the "root" user account**
- ↘ **Security groups should not allow unrestricted access to Telnet (port 23)**
- ↘ **Enable immutability for Recovery Services vaults**
- ↘ **Relational Database Service (RDS) should not have a Public Interface**

FortiGuard SecOps action box: cloud intrusion (CNAPP/identity-driven exploitation)

CNAPP Hard Stop:
Identity anomaly + discovery APIs = active intrusion. Contain immediately.
Do not wait for malware. Do not wait for impact.

In an industrialized cloud threat landscape, intrusions execute through identity, APIs, and automation. In the cloud, valid credentials are the exploit, and APIs are the execution engine.

Compromise is stopped only by revoking trust faster than attackers can operationalize it.

SOC: Signals · Correlations · Detections · Automation

Phase	SecOps-Ready Actions
Signals to Monitor	▪ Authentication from new IP/ASN/geolocation: treat as trust deviation ▪ Impossible travel events: elevate to identity risk ▪ Auth success after repeated failures/MFA challenges: credential pressure → conversion ▪ Credential-validation APIs (e.g., GetCallerIdentity) from anomalous context: proof of credential testing ▪ Discovery-heavy APIs across services/regions: automation signal (mapping the environment) ▪ Privilege/persistence APIs (role trust changes, new access keys, wildcard policies): trust modification = escalation ▪ Early monetization (SES abuse, sudden compute scale-up, cryptomining patterns): impact staging. [Detection: FortiCNAPP]
Correlation Logic	▪ Identity anomaly → credential validation → discovery APIs (compressed timeframes): chain into "active intrusion." ▪ Identity anomaly + discovery + privilege APIs (overlapping phases): automation concurrency. ▪ Same identity reused across services/regions faster than human workflows: non-human tempo. [Correlation/Case execution: FortiSOC / FortiCNAPP]
Automation & Response	▪ Correlate identity + API behavior + posture + workload context: single CNAPP intrusion case ▪ On High/Critical: suspend credentials / invalidate sessions: revoke trust immediately ▪ Restrict IAM scope + revoke trust relationships (keys, roles, policies): contain privilege expansion ▪ Auto-tag as "Confirmed CNAPP Cloud Intrusion": standardize escalation posture [Control: FortiCNAPP; Orchestration/Case execution: FortiSOC]

Mandate: If identity anomalies and discovery APIs coexist, contain now.

DFIR: Evidence · Reconstruction · Proof of Execution

Area	DFIR-Ready Guidance
Priority Evidence	▪ Cloud auth + activity logs (who/where/when) ▪ IAM policy and trust diffs (what changed) ▪ Access-key/token issuance + session reuse (how trust was minted) ▪ API timeline: validation → discovery → privilege → action (how control was achieved) [Evidence source: FortiCNAPP]
Evidence Chain	▪ Identity anomaly → credential validation → automated discovery → privilege expansion: prove execution, not posture [Case/timeline correlation: FortiSOC, FortiCNAPP]
Expected Behaviors	▪ No malware required; legitimate APIs abused at scale; overlapping phases indicate automation. [FortiCNAPP]
Reconstruction Objective	▪ Prove active credential exploitation; show compressed timelines; attribute impact to industrialized cloud intrusion, not posture alone. [FortiSOC + FortiCNAPP]

CISO: Risk · Exposure · Priorities · Decisions

Dimensions	Decision-Grade Insight
Primary Risk	Attackers convert trust into control faster than teams revoke it: cloud intrusion is a trust-velocity problem.
Exposure Context	Identity sprawl, long-lived credentials, over-permissive IAM, and misconfigurations expand blast radius.
CTEM Priorities	▪ Validate identity trust paths continuously. ▪ Detect discovery-driven behavior early. ▪ Minimize time from anomaly to trust revocation. [CTEM/Detection: FortiCNAPP; Operational execution: FortiSOC]
Strategic Decisions	▪ Treat identity compromise as breach-adjacent. ▪ Authorize early, aggressive identity containment. ▪ Accept controlled disruption to prevent monetization.

Impact: How Industrialized Cybercrime Converts Capability into Damage (FortiRecon Intelligence)

Impact is the product of industrialized operations.

> In industrialized cybercrime, every newly exposed vulnerability is effectively pre-priced for exploitation."

In 2025, impact is no longer an emergent consequence of intrusion. It is the intended output of an industrialized adversary ecosystem. **FortiRecon intelligence** describes modern cybercrime as a multi-lane production system in which access is harvested, exploits are weaponized, execution is automated, and impact is monetized repeatedly and predictably at scale.

Operationally, the FortiRecon multi-lane model can be read as **exploit conversion → control-plane amplification → monetization or persistence → disruption**.

Exploitation supply chain: vulnerabilities as industrial inputs

FortiRecon telemetry shows adversaries actively exploited **635 distinct vulnerabilities** in 2025. The impact signal is not just CVE volume. It is how efficiently vulnerabilities are converted into repeatable intrusion paths that can be reused across outcomes. The most consequential industrial shift is the collapse of the separation between adversary classes.

Exploited vulnerabilities now function as a shared supply chain that can drive both **financial extortion (ransomware)** and **strategic access (APT)** through the same defensive gaps, making risk systemic rather than actor-specific.

Exploit availability is the time-compression engine of industrialized exploitation. A large portion of exploited vulnerabilities were already operationally weaponized: **53.86%** had public PoC exploit code, **31.18%** had fully working exploit code, and **24.88%** had both PoC and working exploits available. The impact is temporal, not theoretical. Publicly available exploits enable automation at scale, allow low-skill actors to achieve high-impact outcomes, and compress the defender's window from days or weeks to same-day windows, **often within 24 hours** or a few days at most. In an industrialized environment, TTE becomes more important than exploit sophistication. Once a vulnerability is weaponized, exposure becomes impact.

"New" does not mean "rare." It means rapidly onboarded. **FortiRecon data** show that of **364 vulnerabilities (57.32%)** that began being exploited in 2025, **212 (58.24%**) were CVEs first observed being used by attackers in 2025. This does not indicate zero-day dominance. It indicates rapid operational onboarding. The operational meaning is that patch lag is assumed, newly disclosed vulnerabilities are treated as immediately exploitable inventory, and **exposure accumulation, not exploit novelty, drives compromise**. In industrialized cybercrime, every newly exposed vulnerability is effectively pre-priced for exploitation.

Attribution does not reduce impact. **43.15%** of exploited vulnerabilities were tied to unknown or unattributed groups, consistent with scale-first exploitation, which increases alert volume and makes targeted intrusions harder to distinguish from background exploitation. Operationally, this creates noise with consequences, including higher alert volume, increased false-positive pressure, and greater difficulty separating targeted intrusion from ambient exploitation.

The blast radius is amplified when exploitation targets control planes. **FortiRecon** notes repeated exploitation impacting products from **Microsoft, Google, Fortinet, Cisco, and Apple**, with the impact shifting toward privilege escalation, lateral movement, and persistence. In industrial-scale attacks, a **control-plane compromise equals an organizational compromise**. Zero-days exist, but the dataset explicitly states they are not the primary driver of impact. Known vulnerabilities, persistent exposure, delayed remediation, and reusable exploit chains continue to affect industrial operations.

Monetization line: ransomware as a continuous, industrial production

If exploitation is the supply chain, ransomware is the monetization line. In 2025, **FortiRecon adversary** telemetry identified **7,831 confirmed ransomware victims** globally during the analyzed period, evidence that ransomware is systematic and continuous rather than episodic. The operating model resembles an industrial market, with **controlled fragmentation** (many groups operating simultaneously) and **concentrated output** (a small subset producing a disproportionate share of victims). Entry barriers are low, but relevance is earned through operational efficiency and sustained tempo.

Table E: Industrial exploitation accelerators (evidence → how impact scales)

Industrial signal	Evidence (FortiRecon)	Operational impact
Shared vulnerability supply chain across adversary classes	APT only, Ransomware only	Risk becomes systemic; one gap can feed extortion and espionage
Exploit packaging availability	PoC: 53.86% · Working: 31.18% · Both: 24.88%	Public exploit availability compresses defensive time; automation-compatible exploitation scales
Rapid operational onboarding	364 (57.32%) began exploitation in 2025	Newly disclosed exposure is treated as immediate inventory
First observed attacker use (base clarified)	212 (58.24% of the 364) first exploited-in-2025 vulnerabilities	"New" is onboarded fast; not zero-day dominance
Long-tail scale without attribution	43.15% unknown/unattributed	Higher alert volume + harder separation of targeted vs background exploitation
Control-plane amplification	Repeatedly impacted: Microsoft, Google, Fortinet, Cisco, Apple	Impact tends to shift toward privilege escalation, lateral movement, and persistence

DFIR decision rule: Unknown attribution (43.15%) signals scale. Elevate priority when exploitation aligns with control-plane platforms, because blast radius shifts from host-level compromise to organizational compromise.

↘ **Key observation:** The ransomware ecosystem is not collapsing or consolidating. It is iterating, with multiple groups achieving operational scale simultaneously.

Impact delivery follows industrial logic. Ransomware targeted environments where disruption yields leverage. Sector distribution reflects economic selection driven by **downtime intolerance** and **operational dependency**, while impact remained broadly distributed across industries and still concentrates where disruption produces the greatest leverage. This breadth means planning cannot be limited to "typical ransomware sectors." **Operational leverage, not sector labels, drives selection**.

Geographically, activity is heavily skewed toward economically mature and highly digitized regions, yet the model scales globally with minimal friction, including sustained activity across LATAM.

Table F: Ransomware production model (scale, concentration, redundancy)

Production signal	Evidence (FortiRecon)	What it means operationally
Systematic victimization	7,831 confirmed victims	Monetization runs continuously, not isolated bursts
Concentration among top producers	Qilin 1,021 · Akira 645 · Safepay 645	Impact concentrates where execution is most efficient
Mature operators beyond the top 3	Play · Incransom · Clop-Leaks · Lynx · RansomHub	Operational maturity, not novelty, drives scale
Brand reuse persists	LockBit 5.0 present (lower victim count not quantified)	Ecosystem iterates; established brands evolve
Long-tail redundancy	New groups appear; examples launch with <10 victims in first active month	Low barriers create a testing long tail; only some reach scale
Emerging strategic signal (not deployed)	PromptLock PoC; not deployed in the wild	Direction of travel; current scale still driven by automation + playbooks

↘

Key observation: Ransomware actors prioritize return on disruption, not proximity. Operations scale globally with minimal geographic constraint.

Operational tempo showed no meaningful downtime across the year. Activity peaked, but the line kept running. Even the lowest months still had hundreds of victims, confirming ongoing campaign execution rather than seasonal spikes.

Table G: Where ransomware impact lands (targets, geography, tempo)

Impact surface	Evidence (FortiRecon)	Why it matters
Top targeted sectors	Manufacturing 1,284 · Business Services 824 · Retail 682 · Construction 601 · Commercial & Residential Construction 356 · Hospitals & Physicians Clinics 285	Targeting tracks disruption leverage and downtime intolerance
Observed breadth includes	Finance · Education · Government · Energy · Transportation · Software	Broad targeting surface extends operational risk beyond “typical” sectors
Geographic concentration	US 3,381 · Canada 374 · Germany 291	High exposure + high payment potential concentrates victims
Consistently targeted Europe	UK · France · Italy · Spain	Mature, highly digitized regions remain recurrent
LATAM sustained activity	Brazil 141 · Mexico 69 · Argentina 50 · Colombia 41	Global scaling includes sustained LATAM operations
Operational tempo (peaks)	March 891 · September 876 · October 900 · November 853 · December 863	Continuous campaign execution with visible peaks

Parallel impact lanes: espionage as a condition, hacktivism as disruption

Not all industrial impact is monetized immediately. **FortiRecon intelligence** identified **250+ distinct espionage-aligned adversaries** active during 2025. Espionage has a cumulative, long-lasting impact. It is designed to persist undetected, and consistently targets **government and public sector, telecommunications and critical infrastructure, and technology, defense, and strategic supply-chain industries**. In an industrialized threat model, espionage impact is not a breach event. It is a continuous condition that manifests as long-term data exfiltration, intellectual property theft, and strategic visibility into national and industrial capabilities.

Hacktivism industrializes a different impact vector of **visibility and disruption. FortiRecon telemetry** across hacktivist channels shows high-volume, message-driven coordination, rapid mobilization around geopolitical events, and a preference for disruption, exposure, and reputational damage over persistence, **via website defacement, data leaks and doxxing, DDoS and service disruption, and psychological or reputational pressure**. Hacktivist impact is measured in visibility and disruption, not dwell time. **Industry targeting varied by campaign, but regional patterns consistently tracked geopolitical tension rather than technical exposure.**

The new breach cost driver is speed. Time-to-exploit is compressing, and blast radius is expanding. If defenders cannot match machine-speed execution, they will continue to inherit machine-speed damage.

FortiGuard SecOps action box: impact of industrialized cybercrime

SOC: Signals · Correlations · Detections · Automation

Phase	SecOps-Ready Actions
Detect	▪ Flag in-the-wild exploitation as an incident signal (regardless of CVSS) using exploit/IPS telemetry on exposed services. [Detection: FortiEDR + FortiGate] ▪ Watch control-plane targeting (identity, network edge, management platforms) where a single compromise creates blast radius. [Detection: FortiGate] ▪ Detect ransomware precursors after exploitation (credential abuse, privilege escalation, encryption staging signals). [Detection: FortiEDR + FortiNDR]
Correlate	▪ Link exploitation ↔ impact workflows (lateral movement / encryption prep) to confirm "exploitation-with-intent," not scanning noise. ▪ Tie exploitation ↔ exploit availability (PoC/working exploit release) to raise urgency when weaponization is public and repeatable. ▪ Detect reuse of the same CVE across attack types (crime + espionage style behaviors) to avoid single-narrative bias. ▪ Track temporal chaining (disclosure → exploit availability → exploitation burst) to anticipate windows of peak risk. [Correlation/Case execution: FortiSOC] [Optional intel enrichment: FortiRecon]
Automate	▪ Treat "exploited vulnerability" as an active incident, not a backlog item, and assign an owner immediately. ▪ Auto-escalate when exploitation targets identity/edge/control-plane services and shows reuse across multiple assets. ▪ Auto-enrich with exploit availability, reuse frequency, and related threat context to drive prioritization. ▪ Apply compensating controls fast (virtual patching / access restriction / segmentation) for control-plane targets while remediation is in flight. ▪ Guardrails to avoid fatigue: do not auto-contain on scanning alone: require exploit execution evidence or chained post-exploitation signals. [Orchestration/Case execution: FortiSOC]

SOC takeaway: If a vulnerability is exploited, time, not severity, is the primary risk factor.

DFIR: Investigate · Attribute · Prove Impact

Area	DFIR Focus
Priority Artifacts	▪ Logs tied to exploited services (VPN/web/identity/management) to anchor time and scope. [Evidence source: FortiGate] ▪ Exploit execution traces (payload patterns, malformed requests, injection indicators) to prove execution. [Evidence source: FortiGate / FortiEDR / FortiNDR] ▪ AuthN/AuthZ changes following exploitation (new sessions, role changes, token/credential events). [Evidence source: FortiEDR] ▪ Lateral movement and persistence artifacts enabled post-exploit. [Evidence source: FortiEDR / FortiNDR]
Attack Reconstruction	▪ Timeline: exploitation → access → persistence → impact (prove "active compromise"). ▪ Identify shared exploit chains reused across incidents (repeatability). [Case/timeline: FortiSOC] [Optional intel enrichment: FortiRecon]
Attribution Context	▪ Infrastructure reuse across attempts; distinguish mass exploitation vs targeted intrusion. ▪ Assess overlap between ransomware-linked and espionage-like paths (don't assume motive).
Impact Validation	▪ Prove how exploited CVEs enabled outcomes (encryption, data access, control-plane compromise). ▪ Distinguish exposure-only from exploitation-with-impact; absence of artifacts can be a signal in industrialized ops.

DFIR takeaway: Exploitation evidence is now part of impact analysis, not just initial access.

CISO: Risk · Exposure · Strategic Decisions

Dimensions	Executive Guidance
Risk Reality	▪ Impact is the intended output of industrialized operations; the same CVEs fuel ransomware, espionage-style persistence, and mass exploitation.
Exposure Priorities	▪ Remediate vulnerabilities already exploited, not theoretical risk; reduce exposure across identity, edge, and management control planes. [Operational execution: FortiSOC; Optional intel enrichment: FortiRecon]
Strategic Shift	▪ Patch SLAs alone are insufficient. Organizations must operate CTEM with exploitation evidence driving priority and ownership.
Investment Focus	▪ Fund automation to reduce exploit-to-containment time and prioritize by exploitability + reuse, not severity alone. ▪ Use telemetry-driven prioritization to validate defensive velocity, not static posture.
Board Narrative	▪ The primary risk is exploitation at machine speed; speed is now a budget decision to fund velocity, not volume.

CISO takeaway: Impact is driven by speed, reuse, and scale, not by attacker identity.

The new breach cost driver is speed. Time-to-exploit is compressing, and blast radius is expanding.

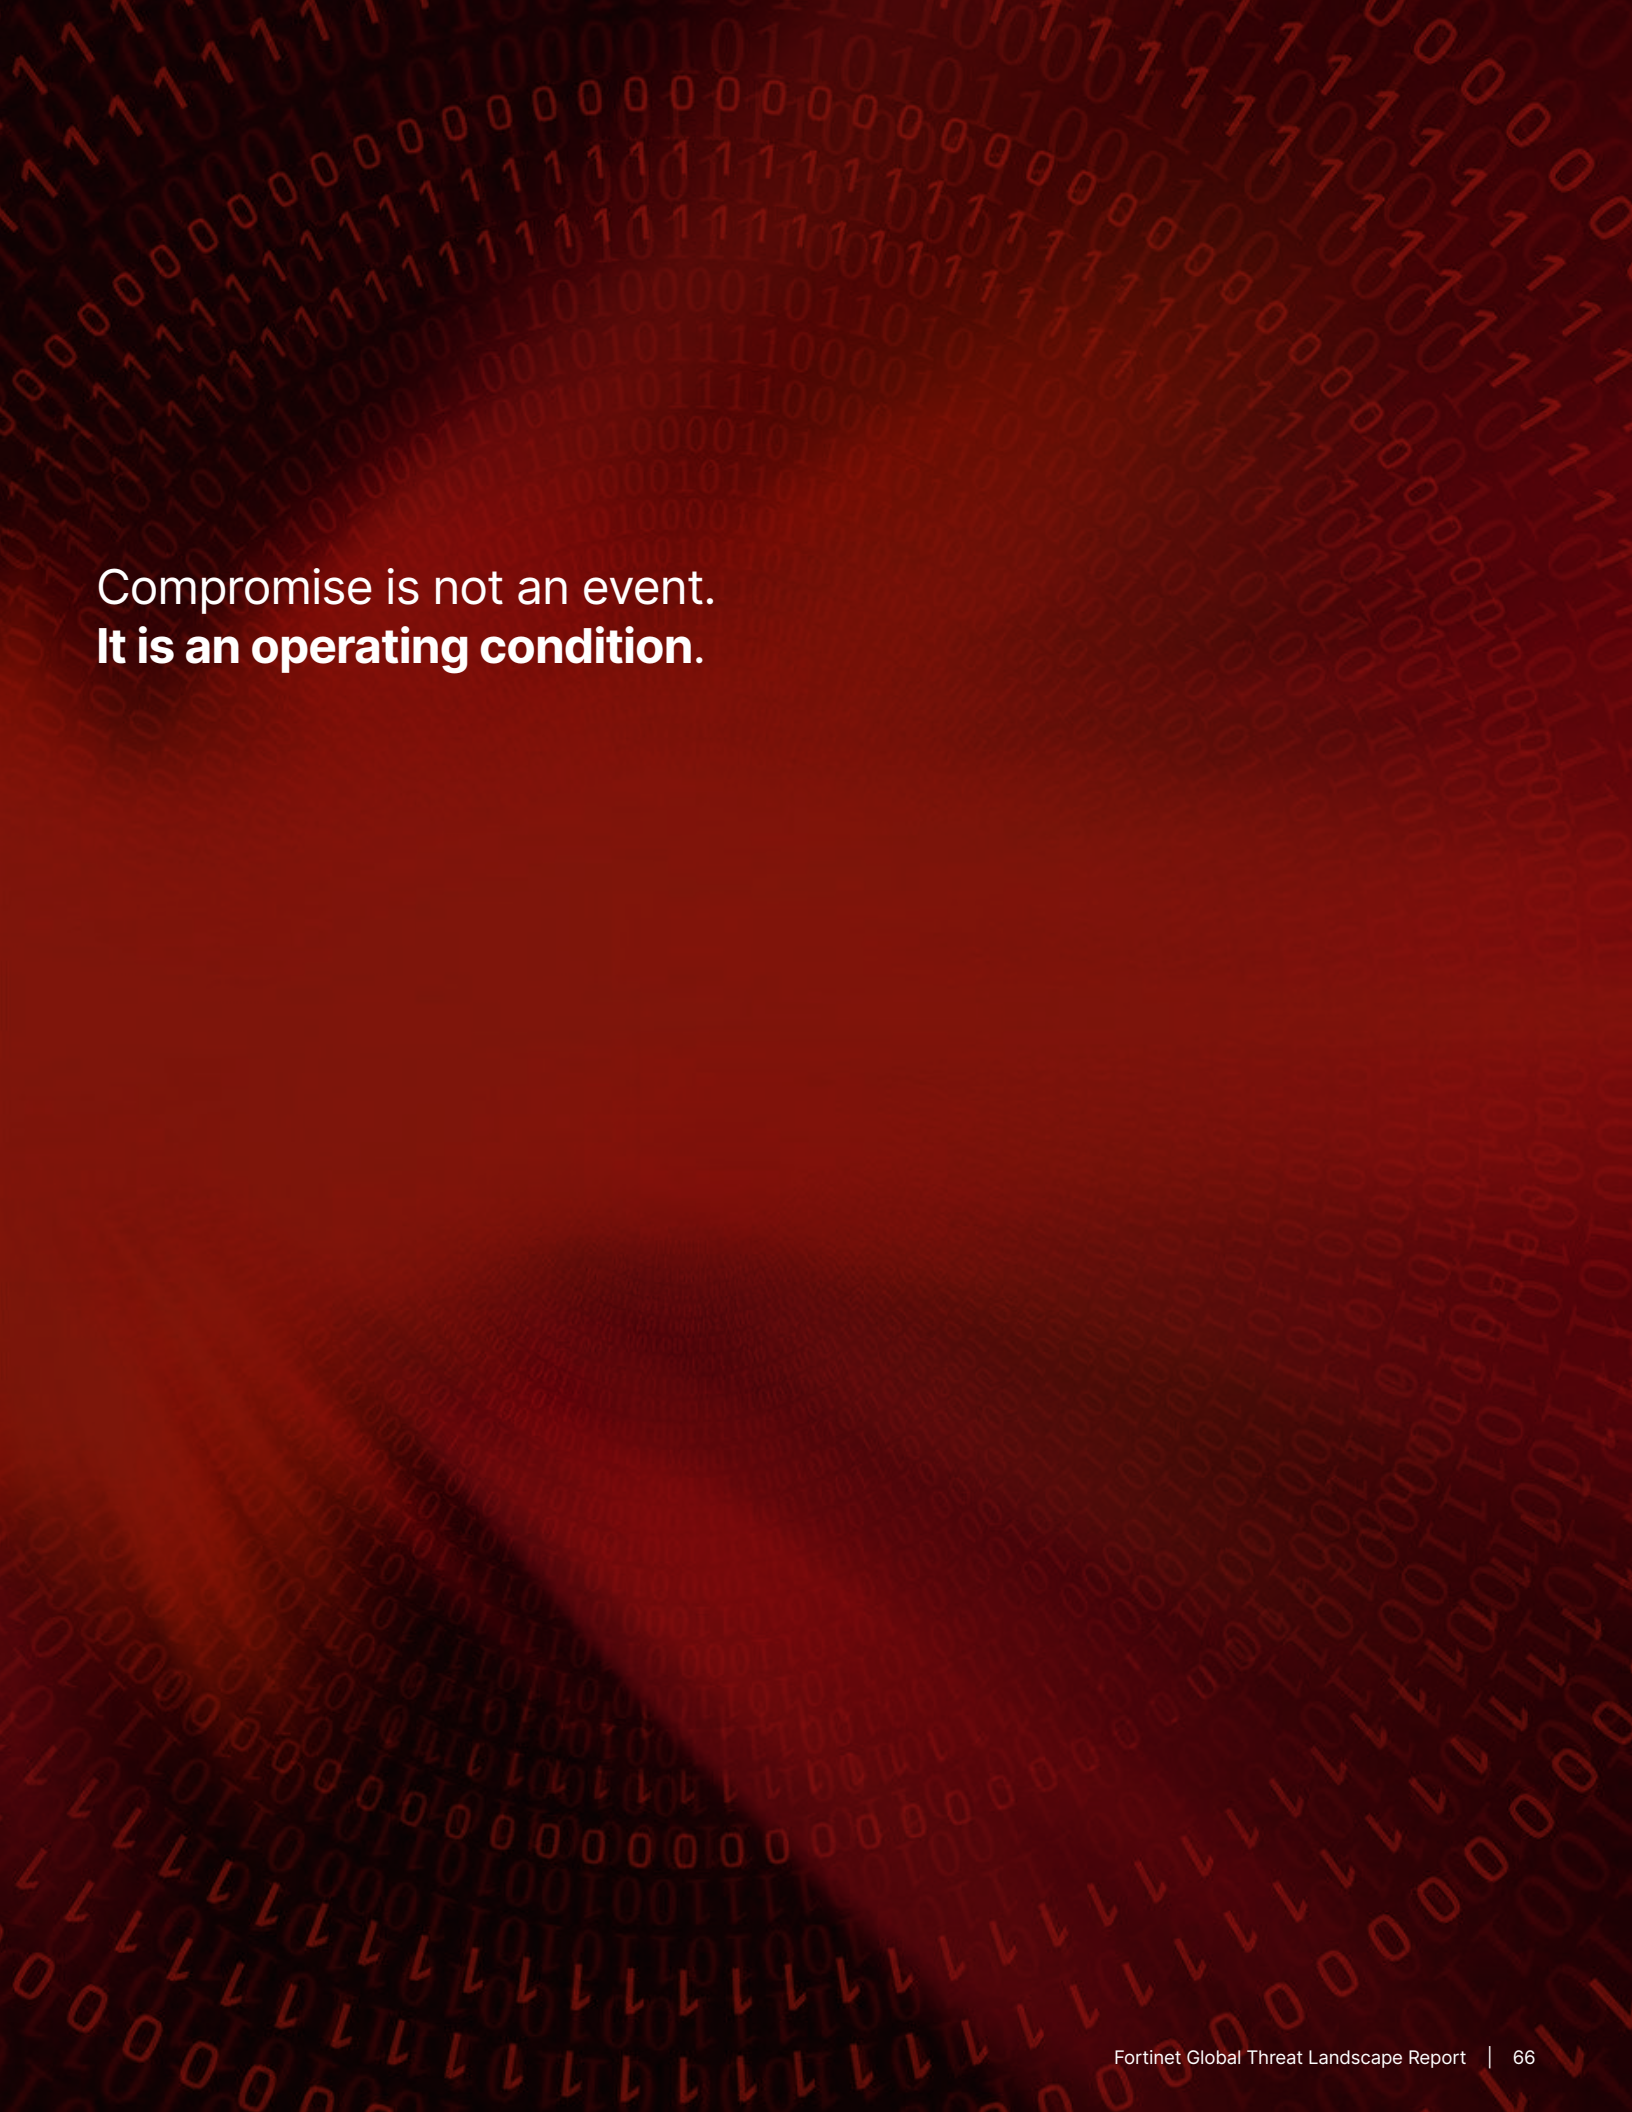
Compromise is not an event.
It is an operating condition.

Conclusion: Restoring Defender Advantage in an Industrial Threat Era

The telemetry FortiGuard Labs gathered over 2025 describes an industrialized system of continuous reconnaissance, sustained credential abuse, automated exploitation, and persistent command-and-control. There are no pauses and no comfortable response windows. Instead, automated pipelines and AI are converting exposure into intrusion at machine speed.

The strategic shift is clear: **Risk is no longer defined by sophistication. It is defined by speed.**

Exploitation now activates when exploit material becomes operationally ready, not when a vulnerability is labeled "critical." Once proof-of-concept or working exploit code is public, exposure becomes a race condition. In an industrial model, **eventual exploitation is the default** unless defenders remove exposure faster than attackers can automate it.

Identity has become the upstream supply chain. Billions of stealer logs and stolen credentials confirm that valid access is continuously harvested, packaged, and reused. **In cloud environments, valid credentials are the exploit, and APIs are the execution engine**. Waiting for malware indicators is waiting too long.

Impact is equally industrialized. Ransomware operates as a steady-state production model rather than as episodic campaigns. Post-exploitation is sustained by a persistent command infrastructure on a global scale. **Compromise is not an event. It is an operating condition**.

Defensive velocity is now the primary control for organizations. Because of this, CTEM cannot be an inventory exercise. It must prioritize what is already exploited, what is automation-ready, and what reduces time-to-containment. If cybercrime runs as an industrial system, defense must likewise industrialize **correlation, automation**, and **containment**, or absorb machine-speed damage as a business cost.

CISO: conclusion table

Strategy · CTEM · Executive Investment

Focus Area	2025 Reality	Fortinet SecOps Leverage	Executive Action
Risk Is Velocity	Exploitation runs continuously at global scale	Unified telemetry across FortiGate, FortiEDR, FortiNDR, FortiCNAPP	Measure and report defensive velocity as a board KPI
Exploited Vulnerabilities Drive Impact	Readiness, not severity, defines risk	FortiGate IPS + KEV/EPSS prioritization + CTEM workflows	Fund exploit-aware remediation, not backlog reduction
Identity Is the Control Plane	Credentials function as the exploit	FortiCNAPP (identity + API correlation), FortiEDR, FortiSOC	Invest in identity-centric detection and trust revocation speed
Post-Exploitation Is Infrastructure	C2 operates persistently at scale	FortiGate Botnet Intelligence + FortiNDR C2 detection	Treat C2 as confirmed compromise, not advisory signal
Industrialized Threat Requires Industrialized Defense	Manual processes cannot match automation	FortiSOC orchestration across domains	Prioritize automation over tool sprawl

SOC: conclusion table

Signals · Correlation · Automation

Operational Focus	2025 Reality	Fortinet Technology Anchor	Required SOC Shift
Pre-Intrusion Chaining	Recon → brute force → exploit overlap	FortiGate IPS + FortiSOC correlation	Escalate chained signals, not single alerts
Exploitation Pressure	Continuous exploitation cadence	FortiGate IPS + FortiEDR execution telemetry	Correlate network exploit with host execution
Fileless / LOLbin Execution	Legitimate tools abused at scale	FortiEDR behavioral detection	Hunt execution patterns, not signatures
Post-Exploitation Control	Persistent beaconing and lateral movement	FortiNDR + FortiGate Botnet Intelligence	Auto-escalate C2 + identity anomalies
Automation Gap	Manual triage loses to machine speed	FortiSOC playbooks across IPS/EDR/NDR/CNAPP	Automate isolation, credential revocation, and gating

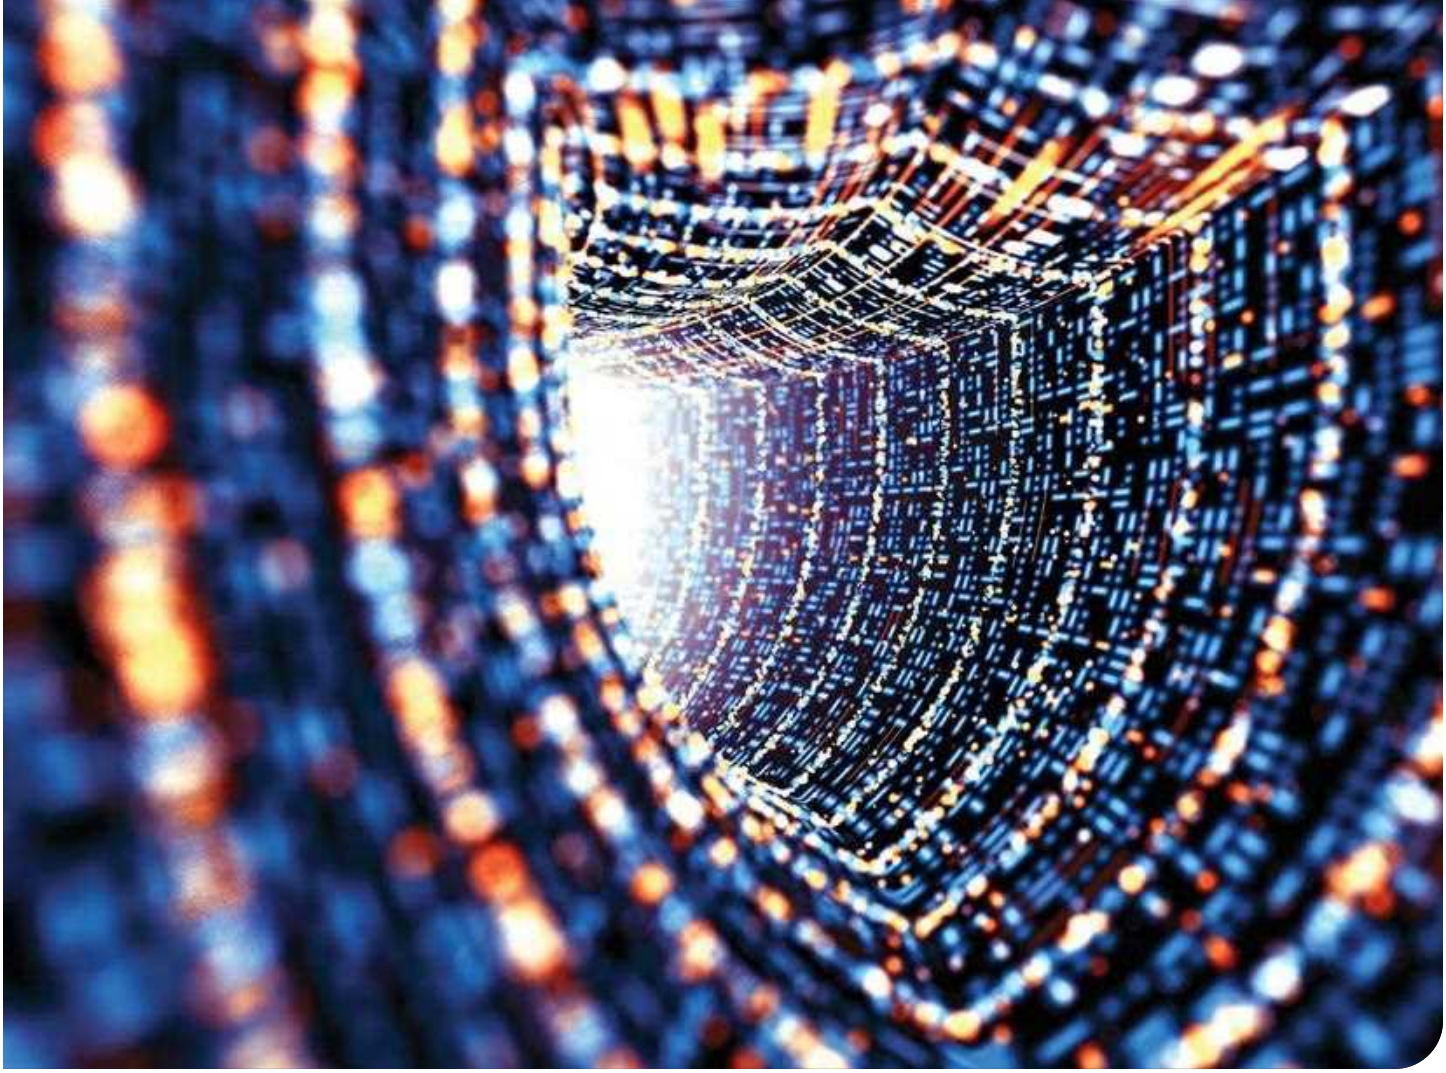

DFIR: conclusion table

Evidence · Reconstruction · Proof of Industrialization

Investigation Focus	2025 Reality	Fortinet Telemetry Source	What DFIR Must Prove
Exploit → Execution Transition	Exploitation converts rapidly to host execution	FortiGate IPS + FortiEDR telemetry	Timeline compression between exploit and execution
Identity-Driven Intrusion	Valid accounts enable scale	FortiCNAPP + FortiSOC + AD telemetry	Credential abuse as intrusion root cause
C2 Confirmation	Persistent command infrastructure	FortiGate Botnet Intelligence + FortiNDR	Active compromise, not attempted intrusion
Automated Discovery & Lateral Movement	RPC, NTLM anomalies, admin tool abuse	FortiNDR east-west visibility	Parallel phases, not linear attack chain
Low Artifact Footprint	Fewer malware artifacts, more native abuse	FortiEDR behavioral signals	Intent and control without relying on malware drops

About the Fortinet Threat Landscape Report

The Fortinet Threat Landscape Report is an annual, data-driven analysis of global cyberthreat activity based on real-world telemetry collected across Fortinet's security platform. It is designed to describe how adversaries actually operate at scale, not how attacks are theorized to occur.

The report draws on global intelligence gathered from millions of network, endpoint, cloud, application, email, and identity control points, as well as FortiRecon dark web intelligence and FortiGuard Labs research. This multi-domain visibility allows the report to track the full operational lifecycle of modern cyberthreats, from exposure and weaponization through exploitation, post-exploitation, and impact.

Rather than focusing on individual malware families or isolated incidents, the report analyzes structural patterns in adversary behavior. Its goal is to identify how cybercrime operates as a system, including how vulnerabilities are operationalized, how access is reused, how automation changes attacker economics, and how speed increasingly determines outcomes.

Each edition is structured around the FortiGuard SecOps chain to align threat intelligence with security operations decisions. The intent is not only to describe what is happening, but to provide defenders with a practical framework for prioritization, response, and exposure management in an environment where threats operate at machine speed.

About FortiGuard Labs

FortiGuard Labs is Fortinet's global threat intelligence and research organization. It is responsible for collecting, analyzing, and operationalizing security telemetry from across Fortinet's installed base, including network, endpoint, cloud, email, and application security platforms.

FortiGuard Labs operates one of the industry's largest threat intelligence ecosystems, combining real-world telemetry, advanced analytics, and direct threat research to track adversary behavior across the full attack lifecycle. This includes vulnerability exploitation, botnet and command-and-control infrastructure, credential theft ecosystems, ransomware operations, and cloud and identity abuse.

The intelligence produced by FortiGuard Labs powers Fortinet's security services, informs global law enforcement and industry partnerships, and supports coordinated disruption efforts such as INTERPOL operations, the Cybercrime Atlas, the Cyber Threat Alliance, and Fortinet's Cybercrime Bounty program.

In addition to detection and response, FortiGuard Labs focuses on reducing adversary capacity by helping dismantle the infrastructure and ecosystems that enable cybercrime at scale.

About Fortinet

Fortinet is a global leader in cybersecurity, driving the convergence of networking and security to secure enterprises, service providers, and governments worldwide. With the industry's largest integrated portfolio of enterprise-grade security solutions, Fortinet protects organizations worldwide across networks, endpoints, cloud environments, applications, and data.

At the core of Fortinet's strategy is the Fortinet Security Fabric, which delivers broad, integrated, and automated protection across the entire digital attack surface. This unique platform approach enables organizations to reduce complexity, improve visibility, and increase response speed in an environment where cyberthreats operate at machine speed.

Fortinet solutions are powered by FortiGuard Labs threat intelligence and supported by a global ecosystem of partners, customers, and public-sector collaborators. Together, Fortinet and FortiGuard Labs work not only to defend organizations but also to disrupt cybercrime operations, increase adversary cost, and help shift the balance back in favor of defenders.

April 29, 2026 2:45 PM / MKTG-3124-0-0-EN

www.fortinet.com

FORTINET