



Cloudflare Threat Report

How adversaries are weaponizing
the Internet

2026

Table of contents

6	Emerging attack techniques and trends
13	Nation-state analysis
33	Cybercrime
46	Community and regional perspectives
54	Recommendations



About Cloudforce One

Driven by a mission to help defend the Internet, [Cloudforce One](#) leverages telemetry from Cloudflare's global network — which protects approximately 20% of the web — to drive threat research and operational response, protecting critical systems for millions of organizations worldwide.

Executive summary

The Internet was built on trust and interoperability, but its architects never anticipated the adversaries who would weaponize those very principles against it.

Every single day, Cloudflare's network is the first line of defense against more than 230 billion threats. Threat actors are moving with the same speed and sophistication as defenders, leaving no room for error in a neck-and-neck battle for the network.

But Cloudflare doesn't just react to these threats; we predict and interdict them. By processing over 20% of the world's Internet traffic, we see the first tremors of an attack before they become a global earthquake and distill patterns across our network to help us anticipate and plan for future disruption. This unmatched visibility allows us to turn raw data into proactive defense.

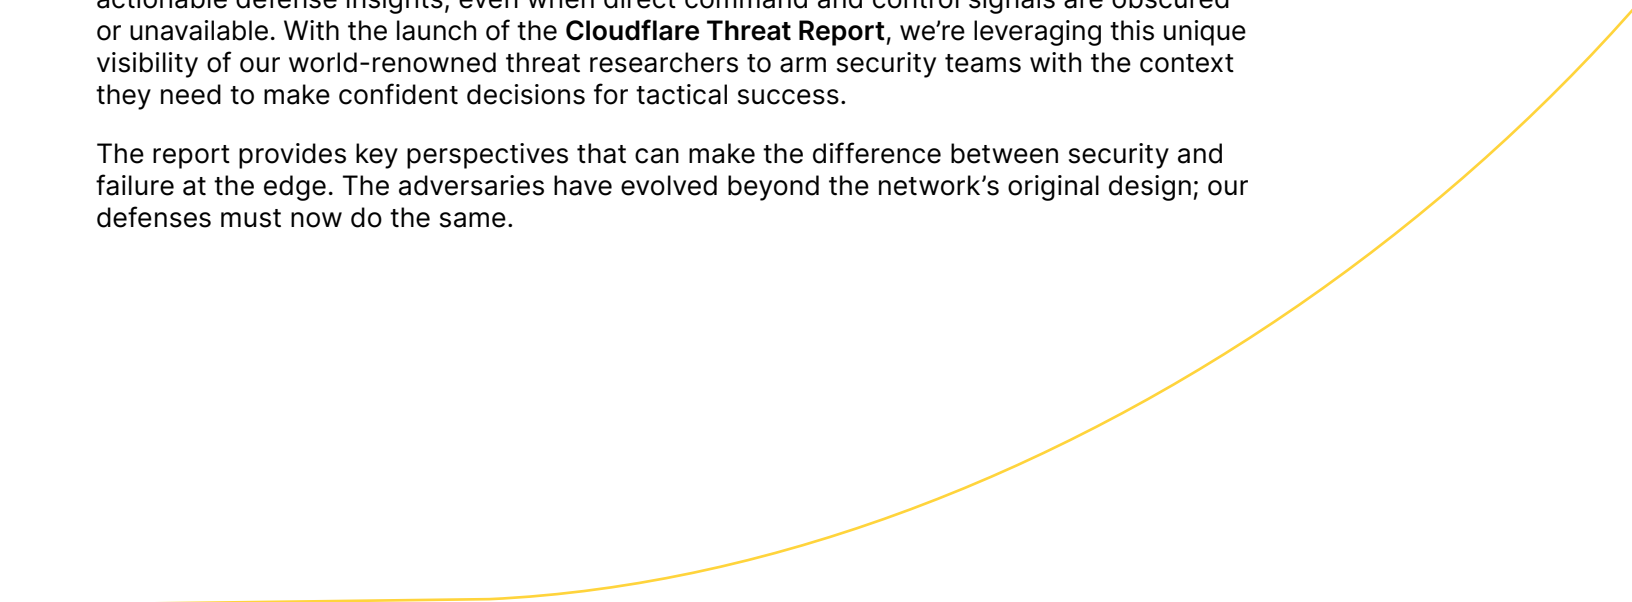
The 2026 threat landscape is poised to reward the stealthy over the loud, and organizations need to know the threats in order to thwart them. To meet this moment, Cloudforce One has spent the last year translating trillions of network signals and threat actor tactics, techniques, and procedures (TTPs) into the insights and recommendations organizations need to prepare for and execute actionable defense. **The inaugural Cloudflare Threat Report is the result of this work.**

Our unmatched global telemetry has made the verdict clear: In 2026, we are witnessing **the total industrialization of cyber threats**, where the barrier to entry has vanished and the "interactive hack" is now a scalable and automated model.

The collapse of the traditional perimeter has turned identity into the primary target, while the explosion of AI and SaaS-to-SaaS complexity has given attackers a force multiplier to move through networks at machine speed. From hyper-volumetric DDoS strikes that paralyze infrastructure to the silent infiltration of corporate payrolls, the 2026 landscape is defined by a shift from brute force to high-trust exploitation.

Our team specializes in turning global signals and critical network interconnections into actionable defense insights, even when direct command and control signals are obscured or unavailable. With the launch of the **Cloudflare Threat Report**, we're leveraging this unique visibility of our world-renowned threat researchers to arm security teams with the context they need to make confident decisions for tactical success.

The report provides key perspectives that can make the difference between security and failure at the edge. The adversaries have evolved beyond the network's original design; our defenses must now do the same.



Key findings for 2026



AI is automating high-velocity attacker operations

The primary metric for risk in 2026 is the Measure of Effectiveness — the ratio of attacker effort to operational outcome. The accessibility of generative AI significantly lowers the barrier to entry for highly effective operations, moving the industry beyond technically elegant code to “offense by the system.” By leveraging a victim’s own cloud, software as a service (SaaS), and AI infrastructure to fund and scale missions, adversaries are achieving a level of frictionless scale that traditional risk models fail to capture.



State-sponsored pre-positioning is compromising critical infrastructure resilience

Chinese threat actors, notably Salt Typhoon and Linen Typhoon, are prioritizing North American telecommunications, government, and IT services for persistent pre-positioning. This strategic targeting suggests a deliberate shift toward preparing for future disruptive events over immediate espionage. By embedding footholds within core infrastructure, adversaries are eroding the foundational resilience of essential public and private sector services, anchoring their presence for long-term geopolitical leverage.



Over-privileged SaaS integrations are expanding the blast radius of attacks

The security of corporate data is now defined by third-party integrations rather than the traditional network perimeter. In 2026, a single over-privileged SaaS-to-SaaS connection can be weaponized via AI to trigger surgical, multi-tenant breaches across entire ecosystems simultaneously. This structural vulnerability turns the “connective tissue” of modern enterprises into a primary vehicle for widespread and automated operational disruption.



Adversaries are subverting service ecosystems to mask attacks

Threat actors are weaponizing legitimate cloud ecosystems (SaaS, IaaS, and PaaS) to camouflage malicious actions within benign enterprise operations. In 2026, the use of trusted platforms for encrypted command delivery has matured into a standardized obfuscation layer within broader, multi-stage hybrid infrastructures. This democratization of scalable, high-bandwidth cloud resources allows even low-tier actors to execute sophisticated attacks that bypass traditional egress filtering.

Key findings for 2026



Deepfake personas are embedding adversarial operatives within Western payrolls

The industrialization of fraudulent identities now allows state-sponsored operatives to embed themselves directly into Western payrolls. These actors leverage deepfake profiles and remote laptop farms to maintain a residency illusion that evades geolocation and identity controls. This infiltration turns the remote workforce into an attack vector, placing malicious insiders within the organization's most trusted administrative and financial systems.



Token theft is neutralizing multi-factor authentication

Adversaries are neutralizing standard multi-factor authentication (MFA) by transitioning from "attacking the box" to "attacking the session." Using infostealers like LummaC2, attackers actively harvest live session tokens to capture already-authenticated states and bypass perimeter controls. This shift has turned ransomware into a simple login event, where attackers exploit fragmented identity estates to move laterally without triggering the credential alerts once relied upon for detection.



Relay blind spots are enabling internal brand spoofing

Attackers are exploiting a critical blind spot where mail servers fail to reverify a sender's identity after a message passes through a third-party gateway. Because the traffic arrives from a "trusted" relay, the system incorrectly treats external spoofed messages as internal or safe. This allows phishing-as-a-service bots to bypass standard protection and deliver high-trust brand impersonations directly to user inboxes by abusing fragmented mail authentication.



Hyper-volumetric strikes are exhausting infrastructure capacity

Hyper-volumetric distributed denial-of-service (DDoS) attacks, fueled by massive botnets like Aisuru, have established a record-breaking 31.4 Tbps baseline that physically exhausts most organizations' network capacity. These autonomous strikes peak in seconds, effectively closing the window for human intervention and placing an extreme resource tax on local infrastructure.

Emerging attack techniques and trends

As we look toward the threat landscape of 2026, the industry must undergo a fundamental shift in how it categorizes risk. For decades, the “sophistication” of an adversary — the technical elegance of their code or the novelty of their zero-day — was the primary barometer for danger. Today, that metric is being replaced by a more pragmatic calculation: the Measure of Effectiveness (MOE).

MOE evaluates a threat by the ratio of attacker effort to operational outcome. In layman's terms, it is a measure of “bang for the buck” — where a high-MOE attack achieves maximum disruption with minimum cost. For example, rather than spending millions to develop a custom exploit, a 2026 adversary might use a low-cost GenAI subscription to automate credential harvesting across thousands of targets. This model measures velocity from initial access to exfiltration and the frictionless scale that attackers leverage to target hundreds of victims at once.

This approach exploits a profound resource asymmetry, repurposing a victim's own cloud, SaaS, and AI infrastructure to fund and execute the mission. This section explores how attackers have moved beyond the pursuit of technical brilliance in favor of industrialized, high-MOE campaigns that turn the connective tissue of the modern enterprise into its primary vulnerability.



MOE in the SaaS supply chain: Weaponizing connective tissue

Cloudforce One has observed an escalation in attacks targeting the connective tissue of the modern enterprise: the integrations between third-party SaaS platforms. Most notably, through our investigation into the threat actor [GRUB1](#), we have seen how a single compromise of a trusted integration — such as the Salesloft Drift connection to Salesforce — can create a dangerous ripple effect, cascading through entire ecosystems to expose hundreds of corporate tenants simultaneously.

In this attack, the adversary bypassed traditional perimeter defenses by targeting high-value credentials buried in code and repository history using automated secret-scanning tools like TruffleHog. Once these keys to the kingdom were harvested, the actor leveraged generative AI in real time to navigate unfamiliar, complex SaaS environments. **Cloudforce One identified the GRUB1 threat actor using AI to pinpoint specific database tables that contained the most valuable information** just moments before gaining unauthorized access to production instances.

The implications of this shift are profound for 2026. The GRUB1 campaign demonstrates that unsophisticated, individual actors can now execute high-impact breaches by industrializing two key phases of the attack:

- **Automated credential discovery:** Moving from manual guessing to automated scanning of the SaaS supply chain.
- **AI-assisted navigation:** Using large language models (LLMs) to bridge knowledge gaps in specialized software like Salesforce, allowing attackers to locate and exfiltrate sensitive data with surgical precision.

This SaaS-to-SaaS pivoting represents a new frontier of risk, where the security of your data is only as strong as the most over-privileged integration in your tech stack.

The GRUB1 campaign demonstrates that unsophisticated, individual actors can now execute high-impact breaches.



MOE in AI: From productivity tool to automated exploit vector

The meteoric rise of AI LLMs has introduced a dual-front risk that blurs the line between a productivity tool and a primary threat vector.

On one front, there is **the unwitting user risk**: The unprecedented adoption rate among consumers and enterprises means that vast quantities of proprietary source code, financial details, and personally identifiable information (PII) are being routinely funneled into these systems. This creates a massive aggregation of sensitive data — a data gravity effect where the AI system itself becomes the most lucrative target for future exfiltration.

In other words, the risk is no longer just a single leaked document, but the potential for a determined adversary to compromise the “corporate brain” through the very tools designed to enhance it.

On the opposing front is **the malicious threat**: the reality that this same technology serves as a force multiplier for attackers.

The historical emphasis on the inherent sophistication of the threat actor or the toolset they employ is becoming increasingly less impactful. An actor who previously lacked the skills to craft a convincing phishing email or write custom malware can now leverage an LLM to generate them rapidly and at scale, significantly lowering the barrier to entry for highly effective operations.

The key implication for 2026 is a fundamental shift in the MOE: AI is not just an additive tool, but the driver of a new paradigm. We predict a move toward “offense by the system” for threat actors, necessitating an equal and opposite shift toward “security by the system” for defenders.

As LLM usage is adopted by a wider array of threat actors — from state-sponsored groups diversifying their toolkits to financially motivated cybercriminals and individual hacktivists — the observed shift toward MOE over sophistication will not only persist but is projected to increase dramatically. One such example is using LLMs to bridge the gap between a bug and a functional exploit by automating semantic mapping.

Instead of needing the rare skill to manually chain disparate flaws, for example, an actor can now use AI to identify how a simple UI oversight — like a URL override — can be coupled with a backend process-spawning API to create a critical vulnerability. The primary threat is no longer the rarity of the skill set, but the velocity of the outcome. The sheer volume of these automated, persistent campaigns matters more than the technical elegance of the code, as the cost to discover and weaponize “weird machines” in a supply chain has been effectively commoditized.

This paradigm shift was made concrete in December 2025, when Cloudflare’s product security team bypassed traditional, high-cost research methods to audit a newly utilized AI tool — OpenCode — uncovering a critical exploit chain that perfectly illustrates this new era of automated effectiveness.

Case study

Deconstructing the OpenCode exploit chain through vulnerability analysis

As part of a broader effort to validate the use of AI for vulnerability research and proactive defense, Cloudflare's product security team put the "MOE over sophistication" theory to the test by conducting a vulnerability analysis of [OpenCode](#).

By leveraging OpenCode itself to perform the critical vulnerability analysis on the tool — a practical "dogfooding" approach — the Cloudflare team proactively identified and disclosed the vulnerability and demonstrated "security by the system."

Vulnerability summary

A malicious website could abuse the server URL override feature of the OpenCode web UI to achieve cross-site scripting (XSS) on `http://localhost:4096`. This, combined with the OpenCode API's `/pty/` endpoints for spawning arbitrary local processes, allowed for remote code execution on the local system. This was possible because:

1. The OpenCode web UI on `localhost:4096` exposed a `/pty/` API endpoint for spawning local processes.
2. The LLM response markdown renderer lacked sanitization, allowing for HTML / JavaScript injection and resulting in XSS on the `localhost:4096` origin.
3. The web UI allowed overriding the server URL via a `?url=` parameter, which could be exploited to load a malicious, attacker-controlled chat session, injecting the XSS payload.

Cloudflare's assessment led the OpenCode developer to fix the vulnerability. This exercise underscores a new paradigm in vulnerability research, showcasing how the same AI-powered tools that accelerate development can be turned inward to enhance security, providing a model for how organizations can utilize these advanced LLM-based solutions to proactively harden their own supply chains.

Alternatively, these vulnerabilities can also be put up for sale on the dark web, vulnerability and exploit marketplaces, or simply exploited directly by a determined adversary to take over vulnerable infrastructure. While mainstream fuzzing required larger computer farms, AI has dramatically changed the cost to obtain similar, if not better, results.

MOE in cloud resources: Living off the XaaS (LotX)

The pervasive adoption of anything-as-a-service (XaaS) — including SaaS, infrastructure-as-a-service (IaaS), and platform-as-a-service (PaaS) — by organizations globally has been mirrored by nearly all threat actors, representing a corresponding shift in tactics.

Cybercriminals, nation-states, and individual hackers now routinely leverage public cloud hyperscaler resources like Amazon Web Services (AWS), Google Cloud Platform (GCP), and SaaS offerings, blending their activities into the massive volume of legitimate cloud traffic. Cloud identities and configurations are also now primary threat vectors. Rather than creating new infrastructure, attackers inhabit environments using stolen credentials from initial access brokers (IABs) or leveraging shell companies to appear legitimate. By exploiting tenant misconfigurations — like overly permissive IAM roles or unsecured cloud storage — they move laterally and operate undetected as legitimate users.

Looking ahead in 2026, the exploitation of trusted platforms will likely evolve from a tactic into a standard operational baseline. Attackers can quickly provision infrastructure for command and control (C2) operations, host phishing campaigns, or stage distributed denial-of-service (DDoS) attacks with high bandwidth. Ultimately, this evolution will transform the threat from one of external intrusion into one of architectural subversion. As a result, defenders must move beyond traditional barrier-based security toward a model that can distinguish between legitimate usage and the weaponization of a tenant's own cloud-native resources.

This democratization of vertically and horizontally scalable computing and networking resources, combined with the systematic compromise of identity trust, effectively lowers the barrier to entry for complex operations and makes it easier for even unsophisticated, lone hackers to execute attacks.

Scalability and elasticity as a weapon

Alternatively and often in tandem, SaaS platforms are also being used by threat actors to host, launch, redirect, or scale attacks. For instance, services like Amazon SES and SendGrid, designed for legitimate bulk email delivery, are frequently exploited to launch sophisticated phishing and malware distribution campaigns.

By routing malicious emails through these established and often allowlisted services, attackers bypass traditional spam filters and gain a veneer of legitimacy, increasing the success rate of their social engineering efforts. Furthermore, attackers utilize cloud storage (like Dropbox, Amazon S3, or Google Cloud Storage) to host malicious payloads for ransomware, trojans, or exploit kits. Even services like GitHub are being abused for this purpose.

This tactic makes detection harder, as the network traffic often resolves to reputable cloud provider or SaaS vendor domains, complicating the task of network security teams seeking to differentiate benign corporate cloud usage from active threats.

IaaS has become the high-performance engine of the modern threat actor's toolkit. Beyond simple hosting, attackers weaponize IaaS by creating rogue virtual machines (VMs) and containers within compromised accounts to establish long-term persistence that is insulated from standard endpoint security. Because these resources are managed by the attacker but billed to the victim, threat actors can engage in "resource hijacking" at an industrial scale.

The most common manifestation of this is hyper-scale cryptomining, where actors exploit Amazon EC2 or Google Compute Engine service quotas to spin up hundreds of high-performance GPU instances within minutes of gaining initial access. Beyond financial theft, these hijacked IaaS nodes serve as high-bandwidth jump boxes for scanning the broader Internet or launching lateral movement attacks against other cloud-native organizations.

In 2026, we anticipate that the abuse of IaaS metadata services will move beyond simple escalation to full-scale tenant takeover. By abusing the internal metadata services of these IaaS providers, attackers can often escalate their privileges from a single vulnerable instance to full administrative control of a cloud tenant, effectively turning the organization's own scalable infrastructure against its entire digital estate.

Key LotX findings across tracked threat actor groups

While the exploitation of cloud resources is an established tradecraft, 2025 investigations highlighted an accelerated maturation in nation-state strategy: Actors are continuing to shift from mere infrastructure abuse toward pervasive LotX. We predict that for 2026, threat actors will attempt to standardize these techniques as a strategic aim for their operational playbooks.

FrumpyToad



Logic-based C2

The tactic: Moving “inside the box” of reputable SaaS logic to evade detection.

FrumpyToad weaponizes Google Calendar for its cloud-to-cloud C2 loop. The malware reads and writes encrypted commands directly into event descriptions, allowing communication with infected hosts without ever connecting to a malicious domain.

PunyToad



Encrypted tunneling

The tactic: Utilizing legitimate developer tools to bypass egress filtering.

PunyToad uses tunneling capabilities and cloud computing to create resilient, living-off-the-cloud architectures. This masks the backend origin IPs and prioritizes long-term persistence on critical edge devices.

NastyShrew



Paste site dead drop resolvers

The tactic: Using public “paste” sites to coordinate shifting infrastructure.

NastyShrew uses services like teletype[.]in and reentry[.]co as dead drop resolvers (DDR). Infected hosts poll these sites to retrieve rotating C2 addresses, ensuring the traffic resolves to a benign-looking domain.

PatheticSlug



PaaS-ing the perimeter

The tactic: Exploiting the “reputation shield” of cloud ecosystems to mask malicious delivery.

PatheticSlug used Google Drive and Dropbox to host XenorAT payloads, leveraging GitHub for covert C2. By utilizing these high-reputation providers, they successfully blend into legitimate enterprise traffic and bypass traditional security filters.

CrustyKrill



SaaS-hosted phishing

The tactic: Blending credential harvesting into common cloud hosting.

CrustyKrill hosts C2 pages on Azure Web Apps (.azurewebsites.net) and uses ONLYOFFICE to host payloads. This grants their operations a veneer of legitimacy, as the network traffic resolves to a reputable cloud provider.

The industrialization of insider threats

The Cloudforce One team continues to analyze the critical evolution in state-sponsored insider threats through mapping the industrialization of North Korean remote IT worker schemes. These operatives infiltrate Western organizations by leveraging fraudulent identities and AI-driven deepfakes to bypass video interviews, ultimately funneling hundreds of millions of dollars in revenue back to the regime.

Operationally, these workers maintain the illusion of domestic residency by using US-based “laptop farms” and facilitators to host corporate hardware, while accessing the devices via RMM software from abroad. To establish further legitimacy, thousands of operatives create comprehensive digital personas on platforms like LinkedIn and GitHub, often “renting” the credentials of complicit US citizens.

Despite these sophisticated tactics, several high-fidelity detection indicators have emerged, including “impossible travel” login alerts, the presence of mouse-jiggling software, and specific video metadata micro-artifacts consistent with real-time deepfake rendering. To neutralize this growing threat, for 2026 Cloudforce One recommends that organizations shift from traditional perimeter defense toward zero trust biometric verification and the strict geofencing of all remote management tools.



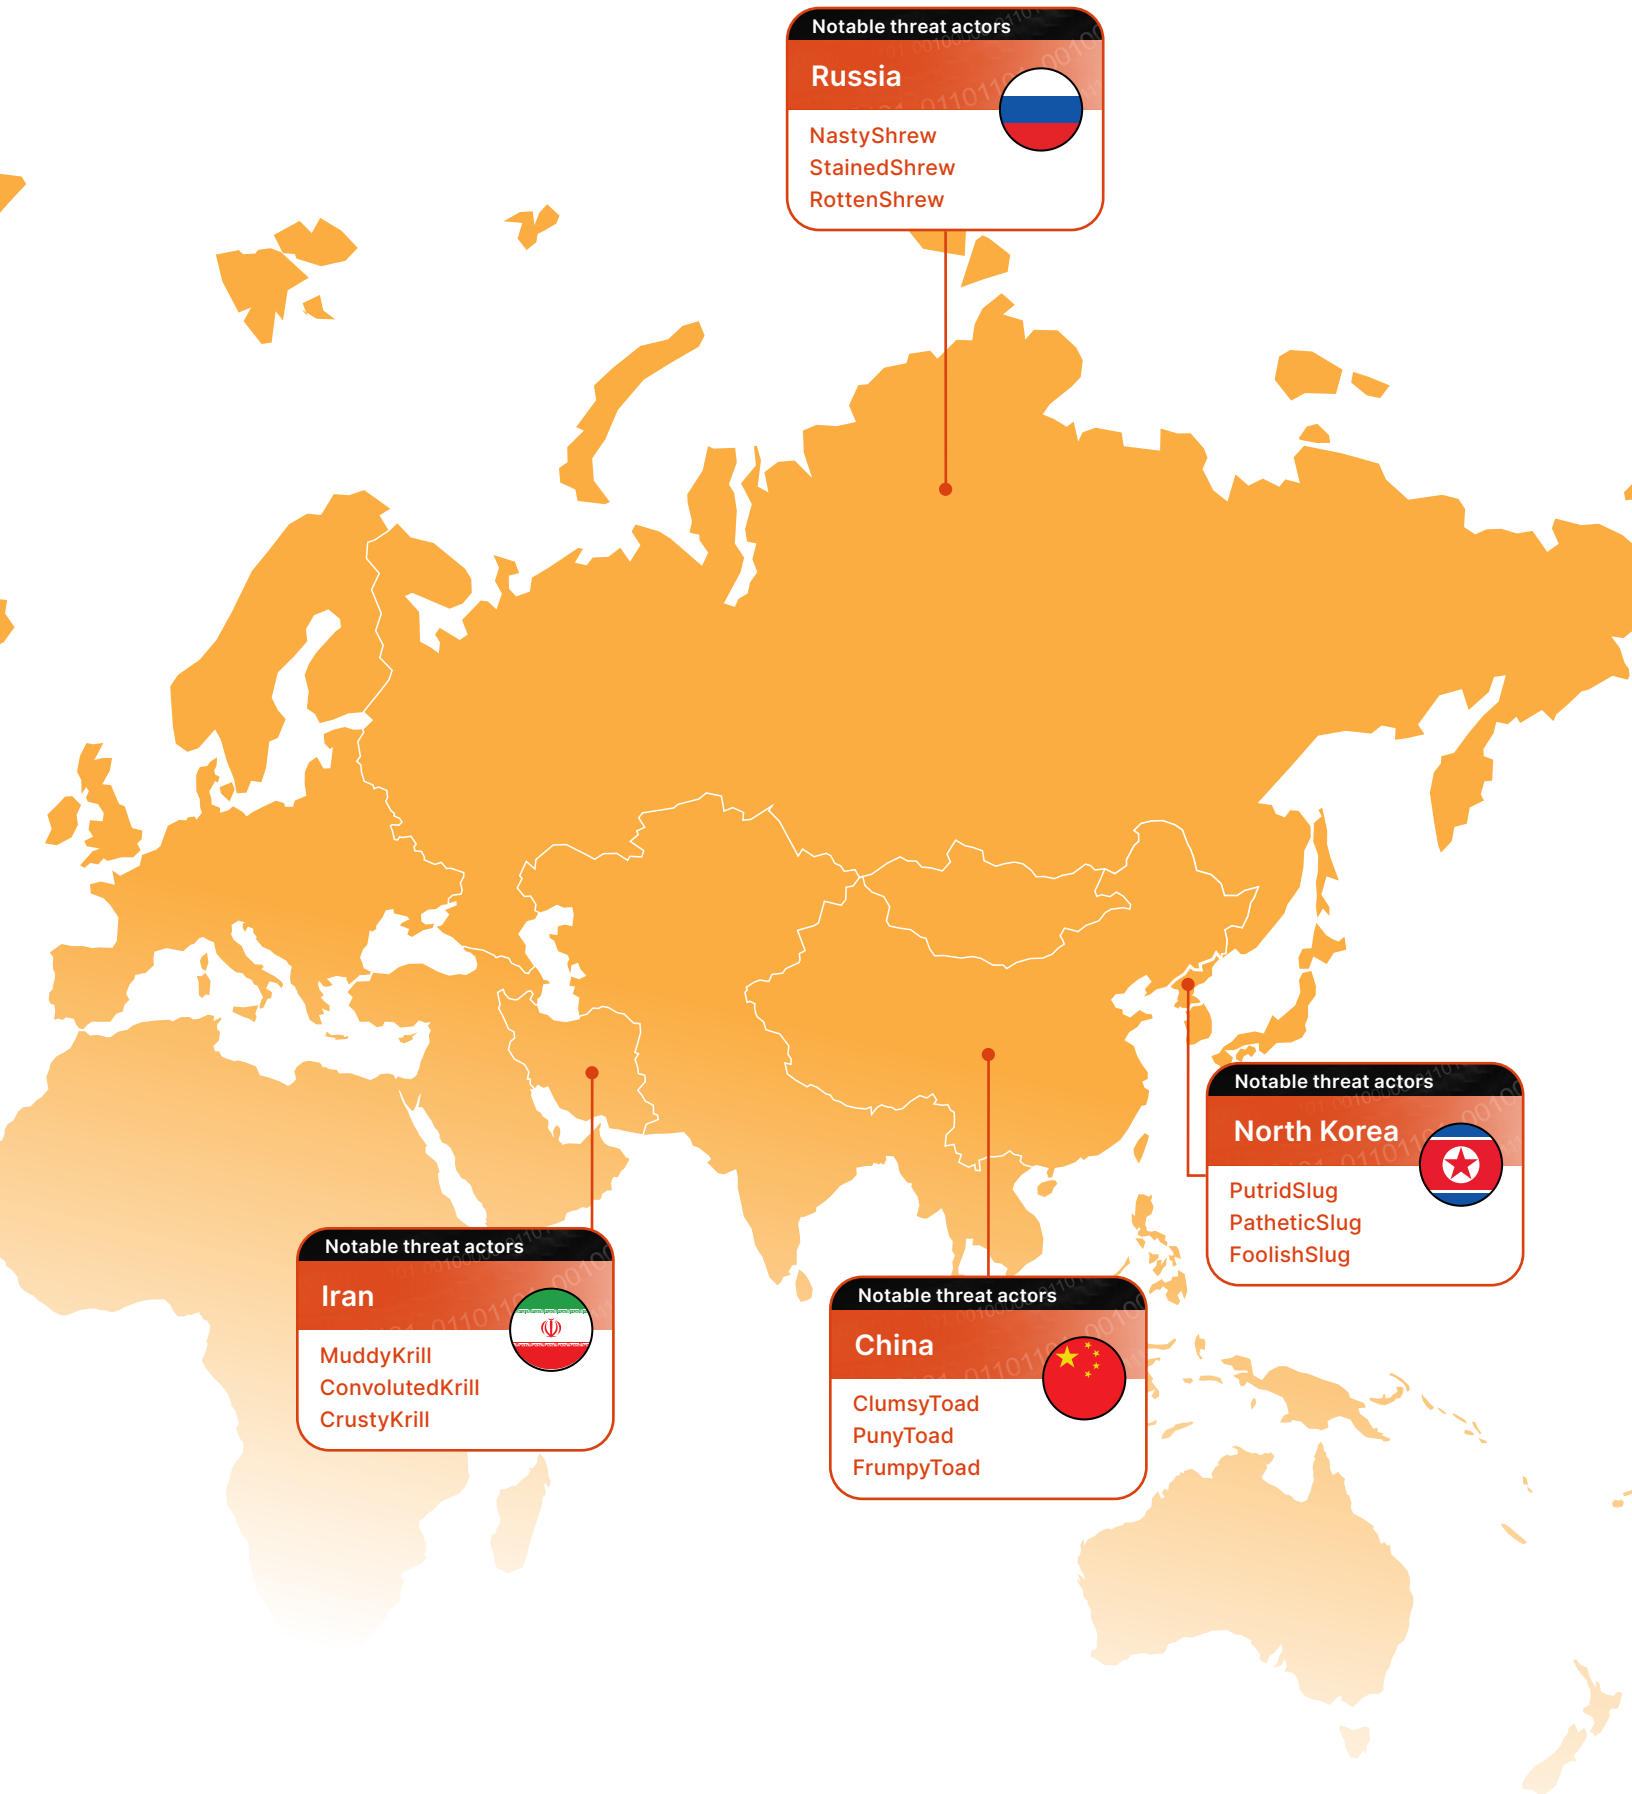


Nation-state analysis

The modern threat landscape from state and state-aligned actors is increasingly defined by the maturation and tactical synchronization of cyber operations with kinetic and geopolitical objectives, a coordination most notably observed in the temporal correlation between disruptive digital strikes and conventional military actions in recent conflicts. Cloudforce One has observed four primary state actors — **Russia, China, North Korea, and Iran** — as they continue to refine a toolkit that extends far beyond simple cyber espionage. As we move through 2026, we expect this trend to continue as state-sponsored cyber operations become an increasingly permanent fixture of physical and political conflict.

This section explores some of the trends, highlights, and examples of approaches used by cyber actors to conduct operations that serve various nation-state strategic goals, ranging from possible support of battlefield operations to the pursuit of financial gains to circumvent sanctions.







Russia: Blurring the line between digital espionage and operational support

The Russian cyber threat continues to operate under a high-frequency, broad targeting model. Groups like NastyShrew maintain their established practice of leveraging high-reputation cloud services to mask C2 infrastructure and employing geofencing to evade global security scanners, sustaining persistent access to Ukrainian critical systems. Russian capabilities are mature and effective and continue to support the geopolitical objectives of the Russian state as exemplified by StainedShrew's ongoing phishing campaigns against NATO allies and RottenShrew's geolocation campaigns against tactical communication apps used by the Ukrainian military, often with temporal correlations to Russian kinetic military operations.

We expect these high-frequency operations to continue through 2026.

Threat actor profile



NastyShrew



Aliases / AKAs: UAC-0010, Gamaredon, Armageddon, Aqua Blizzard, Primitive Bear, UNC530, ACTINIUM, IRON TILDEN, Shuckworm, DEV-0157, Trident Ursa



Primary objectives

Support of battlefield operations and persistent access to critical systems



Targeting

Ukrainian government and critical infrastructure



Tools and TTPs

Pteranodon family, VBScript chains, dead drop resolvers (paste sites), and geofencing

Strategic profile and context

NastyShrew is focused on high-frequency persistence and the disruption of Ukrainian critical infrastructure.

NastyShrew (often tracked as Gamaredon or Primitive Bear) remains one of the most active and persistent Russian threats tracked by Cloudforce One. Their operations are characterized by a high frequency of campaigns casting a wide net and maintaining persistence on thousands of endpoints across Ukraine.

Infrastructure and C2 (2025)

NastyShrew has mastered the use of legitimate services to bypass reputation-based filtering.

- **Dead drop resolver:** The group uses established paste sites (e.g., teletype[.]in, reentry[.]co) to host the names of their active C2 tunnels. Infected hosts poll these sites to retrieve the current C2 address, allowing the actors to rotate backend infrastructure in minutes.
- **Geofencing:** Cloudforce One has observed NastyShrew purchasing VPS infrastructure using cryptocurrency. These servers are configured with strict firewall rules that only permit inbound traffic from Ukrainian IP ranges, which complicates research and detection.
- **Attribution:** NastyShrew consistently uses a mix of anonymity services to access their infrastructure. In February and November 2025, several members of this cluster were arrested in Thailand (Phuket) in a joint operation involving US and Thai authorities.¹
- **Delivery:** Relies on high-frequency phishing utilizing .lnk and .ps1 files to trigger multi-stage VBScript chains.

Threat actor profile



StainedShrew



Aliases / AKAs: VoidBlizzard, Laundry Bear, UAC-0190



Primary objective

Diplomatic espionage and strategic opportunism regarding geopolitical milestones



Targeting

Government, defense, transportation, media, NGOs, and healthcare sectors in Europe and North America



Tools and TTPs

Software and security vendor impersonation, PLUGGYAPE backdoor, and high-frequency infrastructure rotation

Strategic profile and context

In May 2025, Dutch intelligence publicly reported on StainedShrew for the first time and attributed the group to Russia. StainedShrew is highly active in targeting events of perceived geopolitical importance.

2025 operational retrospective

Throughout 2025, StainedShrew demonstrated a high degree of strategic opportunism, aligning their operations with major geopolitical milestones and defense-sector events. Their efforts primarily targeted NATO and its member states, with a significant pivot toward Ukrainian interests in the final quarter of the year.

- **Event-driven tactical targeting:** StainedShrew exhibited a persistent focus on high-stakes diplomatic and security conferences to facilitate initial access. This included various high-level meetings such as the NATO Summit in the Hague, the EU Defence Minister meeting, and the Brussels Indo-Pacific Forum. Other targeting of security-related events included the International Defence Exhibition SIDEC in Slovenia and the EU Defense Night in Washington, D.C. Beyond purely diplomatic targets, the group also targeted industry-specific events like the World Agriculture Forum, indicating a broad collection mandate across both political and industrial sectors.
- **Software brand impersonation:** To facilitate credential theft and payload delivery, the actor utilized malicious domains that mimicked common enterprise IT and security products. By spoofing portals for Microsoft productivity suites, ESET security solutions, and other common antivirus software, StainedShrew leveraged the inherent trust associated with these brands to lower target suspicion.
- **Infrastructure setup:** StainedShrew purchases domains and hosting from a variety of providers. The group anticipates having their infrastructure disrupted and purchases new domains and hosting quickly. StainedShrew's tactics are not unique or especially sophisticated and don't necessarily need to be to accomplish the group's objectives.

Threat actor profile

RottenShrew





Aliases / AKAs: UAC-0185, Lost Potential, Group M, UNC4221



Primary objective

Reconnaissance and espionage possibly in support of kinetic operations



Targeting

Ukrainian military and defense organizations



Tools and TTPs

PINPOINT (Geolocation payload), MESHAGENT (RMM), and Signal Phish Kit

Strategic profile and context

RottenShrew functions as a specialized reconnaissance unit. Their primary goal is the geolocation of military personnel likely to be followed by kinetic action.

2025 tactical focus

- **The “Signal” vector:** RottenShrew conducted a widespread campaign targeting Signal accounts used by Ukrainian military personnel mimicking the *Kropyva* application (a proprietary artillery guidance system used by the Armed Forces of Ukraine) to entice users to link their Signal accounts to RottenShrew-controlled instances, giving the actors full access to communications and metadata.
- **Targeting Ukraine:** Besides the *Kropyva* application, RottenShrew targeted the military application Delta, Teneta-related sites (a tech company providing tactical modem and blue force tracking devices), the digital government services application Diia, and the Ukrainian border crossing service e-Cherha.
- **Operational objectives:** RottenShrew conducts espionage campaigns, harvesting credentials and collecting information from victim devices. Furthermore, the group geolocates its victims — a form of reconnaissance indicating a potential interest and overlap with kinetic targeting.

Tool	Function
PINPOINT	A lightweight JS payload using the browser’s Geolocation API to extract high-accuracy coordinates of victims
MESHAGENT	Legitimate remote management tool used for persistence and live monitoring of victim screens
Signal Phish Kit	Custom framework for device linking and credential theft



China: A sophisticated model of infrastructure pre-positioning and cloud-native stealth

China's cyber strategy continues to evolve beyond traditional bulk data theft to a sophisticated model of infrastructure pre-positioning and cloud-native stealth. While groups like ClumsyToad maintain high-frequency regional espionage using localized lures and USB-borne worms like SnakeDisk, the broader apparatus increasingly prioritizes the long-term compromise of critical edge devices and telecommunications backbones.



In addition, Chinese actors (notably Salt Typhoon and Linen Typhoon) are increasingly prioritizing North American telecommunications (e.g., the breach of AT&T, Verizon, and Lumen), government (e.g., US House of Representatives), and IT services (e.g., the July 2025 Microsoft SharePoint compromise).² These operations represent an expansion of traditional espionage, leveraging deep access to critical sectors to achieve simultaneous intelligence collection and persistent pre-positioning within US infrastructure.

By weaponizing legitimate enterprise ecosystems — such as FrumpyToad's use of Google Calendar for C2 or PunyToad's exploitation of F5 and VMware vCenter and ESXi — actors have created a resilient, living-off-the-cloud architecture that allows for rapid data exfiltration while remaining nearly invisible to standard perimeter defenses.³

We expect this strategy of pre-positioning within telecommunications and critical cloud services to continue through 2026, serving as a permanent foundation for both long-term espionage and potential disruptive operations.

Threat actor profile

ClumsyToad



Aliases / AKAs: BASIN, BRONZE PRESIDENT, Earth Preta, HoneyMyte, LuminousMoth, MUSTANG PANDA, Polaris, Red Lich, Stately Taurus, TA416, TANTALUM, TEMP. HEX, Twill Typhoon

Primary objective
High-frequency regional espionage and long-term compromise of diplomatic entities

Targeting
Southeast Asian governments and European diplomatic entities

Tools and TTPs
Spear-phishing, PlugX / Korplug, Toneshell, Yokai and PUBLOAD, SnakeDisk (USB worm), and DLL side-loading

Strategic profile and context

In 2025, ClumsyToad moved away from standard .lnk files in favor of Windows Management Console (.msc) files. This tactic exploits the inherent trust in system-signed management files to bypass EDR LotL detections.

- **SnakeDisk / Thailand campaign (Sept 2025):** Cloudforce One identified a novel USB worm, SnakeDisk, geofenced to execute exclusively on Thailand-based IP addresses. It propagates via infected removable drives to deploy the Yokai backdoor, specifically targeting Thai government and police infrastructure.

Operational spotlight


ClumsyToad disruption

While Chinese actors continue to evolve their stealth capabilities, 2025 saw a setback for ClumsyToad as Cloudforce One conducted a disruption of the infrastructure supporting their ongoing operations. Leveraging a unique vantage point across the network edge allowed for targeted interventions that effectively neutralized key actor infrastructure and capabilities:


- **Account termination:** 20+ malicious Cloudflare accounts were terminated, severing the actor's ability to manage cloud resources.
- **Edge enforcement:** 400+ domains were blocked at the network edge, neutralizing delivery and phishing redirects.
- **Asset neutralization:** 87 core assets were shuttered, including 70 C2 domains, 10 serverless Workers used for proxying traffic, and 7 Pages used for payload delivery.

This systematic takedown demonstrates that while Chinese actors prioritize cloud-native stealth, their reliance on legitimate enterprise infrastructure creates single points of failure that can be exploited by defenders with sufficient network visibility.


Threat actor profile



PunyToad




Aliases / AKAs: UNC5221, UTA0178, Warp Panda




Primary objective

Exploitation of edge network appliances for persistent pre-positioning



Targeting

Edge network appliances (F5 BIG-IP, VMware vCenter and ESXi)



Tools and TTPs

BRICKSTORM (Go-based ELF), GarbTerm, WebSockets (wss://) C2, and DNS over HTTP (DoH)


Strategic profile and context

PunyToad is a group specializing in the exploitation of edge network appliances for persistent pre-positioning and long-term espionage.


2025 tactical focus

- **F5 BIG-IP breach (Oct 2025):** Cloudforce One confirmed PunyToad maintained persistent access to F5 systems for over a year, exfiltrating BIG-IP source code and documentation on undisclosed vulnerabilities.
- **Malware:** Use of BRICKSTORM, a Go-based ELF binary, often obfuscated with Garble.⁴
 - **Reverse logic:** The malware acts as a web server; the C2 connects to it as a client via WebSockets (wss://).
 - **Stealth:** Masquerades as legitimate VMware / Postgres processes (e.g., pg-update, rpclistener) in directories like /opt/vmware/vpostgres/.
- **Infrastructure:** Extensive use of proxying and tunneling techniques. The group utilizes DNS over HTTP (DoH) to resolve C2 domains, which proxy traffic through tunnels to bypass traditional firewall rules.


Threat actor profile



FrumpyToad




Aliases / AKAs: APT41, Amoeba, BARIUM, BRONZE ATLAS, BRONZE EXPORT, Blackfly, Brass Typhoon, Earth Baku, G0044, G0096, Grayfly, HOODOO, LEAD, Red Kelpie, TA415, WICKED PANDA, WICKED SPIDER




Primary objective

Global financial espionage and LotX stealth operations



Targeting

Global financial sector and SaaS providers



Tools and TTPs

TOUGHPROGRESS, KEYPLUG, ShadowPad, DUSTTRAP, DUSTPAN, and Google Calendar-based C2 mechanisms

Strategic profile and context

FrumpyToad is a group that utilizes LotX tactics to blend C2 traffic into legitimate enterprise ecosystems.

2025 tactical focus

- **The TOUGHPROGRESS campaign (Oct 2024–mid 2025):** Cloudforce One telemetry identified a sophisticated cloud-to-cloud C2 loop where FrumpyToad weaponized Google Calendar to blend in with enterprise traffic.
- **Delivery:** Spear-phishing emails redirect through a chain of serverless platforms to host malicious ZIP archives on compromised government sites.
- **C2 mechanism:** The TOUGHPROGRESS malware reads and writes encrypted commands directly into Google Calendar event descriptions.
 - Stolen data is embedded in event descriptions dated May 30, 2023 (hardcoded).
 - Operators place commands in events dated July 30–31, 2023, which the malware polls and executes.
- **Result:** This allows the actor to communicate with infected hosts without ever connecting to a malicious domain, as all traffic remains within the encrypted ecosystem.

Summary of defensive indicators

Actor cluster	Detection priority	Primary malware
ClumsyToad	Monitor .msc file execution and USB mounting	PubLoad, Yokai
PunyToad	Audit outbound wss:// connections	BRICKSTORM
FrumpyToad	Unusual volume of API calls to Google Calendar	TOUGHPROGRESS, DUSTTRAP



North Korea: Identity-driven infiltration and analytical parasitism

In 2025, North Korea continued to double down on its established strategy of human-centric operations, augmenting traditional infrastructure exploitation with a more aggressive and industrialized weaponization of identity and trust. Driven by the regime's dual requirements for strategic intelligence and illicit revenue, groups such as PutridSlug, PatheticSlug, and FoolishSlug have industrialized the use of generative AI, deepfakes, and analytical parasitism as vectors for seeking unauthorized access to sensitive data and secure environments.

We expect this focus on human-centric exploitation to continue throughout 2026, with the regime further refining these industrialized identity-theft tactics to bypass traditional security filters and obtain consistent financial flows.

Threat actor profile

PutridSlug



Aliases / AKAs: BlueNoroff, APT38

Primary objective
Illicit revenue generation for the North Korean regime by conducting high-stakes financial cybercrime

Targeting
Cryptocurrency, blockchain, and fintech industries

Tools and TTPs
Deepfake video / audio impersonation, trojanized coding challenges, and Smart Contract manipulation

Strategic profile and context

Throughout 2025, PutridSlug augmented its high-volume phishing operations with social engineering powered by generative AI. Their objective was to more effectively circumvent the identity perimeter of financial institutions and crypto-native entities.

- **Deepfake deception:** A hallmark of their campaigns was the use of deepfake video and audio to impersonate company executives during Zoom calls to target tech firm employees. By leveraging the victim's established trust in their leadership, PutridSlug bypassed social defenses to trick employees into downloading malicious payloads.
- **Trojanized coding challenges:** PutridSlug continued to improve and diversify its delivery vectors by embedding tailored malicious code within GitHub-hosted coding assessments targeting prospective job candidates. Once a candidate clones the repository, hidden malicious code specifically crafted for the target's environment is deployed, providing the actor initial access that can lead to financial theft and other devastating breaches.
- **Exploiting smart contract logic:** By compromising the workstations of developers at crypto firms, PutridSlug gained the ability to manipulate smart contract code before deployment, allowing for multimillion-dollar logical thefts rather than just wallet-draining.

Threat actor profile

PatheticSlug



Aliases / AKAs: Kimsuky, Velvet Chollima, Emerald Sleet, APT43

Primary objective

Large-scale cyber espionage to support the North Korean regime's strategic, diplomatic, and nuclear ambitions

Targeting

Government agencies, think tanks, academic institutions, and media organizations

Tools and TTPs

Advanced social engineering and AI integration, browser-centric exploitation, quishing, blend of commodity and custom malware, and LotX mastery

Strategic profile and context

PatheticSlug's 2025 strategy focused on analytical parasitism — attempting to insert themselves into the global policy-making cycle to steal pre-published research and influence diplomatic narratives.

- **The “journalist” persona:** PatheticSlug operators posed as journalists from legitimate news outlets, in attempts to conduct interviews with policy experts to gather off-the-record insights to provide the regime with critical visibility into the diplomatic and military strategies of perceived global and regional adversaries.
- **Contextual account takeover:** They utilized specialized browser extensions designed to mimic legitimate productivity tools. Once installed, these extensions are designed to silently monitor webmail sessions and exfiltrate stolen email addresses, usernames, passwords, cookies, and browser screenshots to actor-controlled infrastructure.
- **Spoofing the “expert circle”:** PatheticSlug targeted global embassies through sophisticated phishing campaigns that leveraged high-fidelity impersonations of trusted diplomatic contacts. By utilizing credible lures, such as official correspondence and event invitations, attackers directed victims to malicious cloud storage links (e.g., Google Drive or Dropbox). These links facilitated the delivery of XenoRAT, a stealthy malware variant that leverages GitHub as a covert C2 channel.

Threat actor profile



FoolishSlug



Aliases / AKAs: Andariel, Silent Chollima, Stonefly, APT45



Primary objective

Self-funding military espionage via ransomware and cryptocurrency theft



Targeting

Cryptocurrency ecosystem; defense, aerospace, nuclear, and healthcare sectors



Tools and TTPs

Ransomware, living-off-the-land (LotL) mastery, cloud tunneling, and edge-device exploitation

Strategic profile and context

FoolishSlug serves as the most tactically aggressive group, frequently utilizing criminal tactics to fund the regime’s high-end military espionage.

- **The espionage-ransomware hybrid:** FoolishSlug refined the double-track intrusion model that synchronizes espionage with financial extortion. By exfiltrating sensitive intellectual property while deploying ransomware variants like Maui or Play, the group maximizes the ROI of the breach. This hybrid approach helps the group achieve fiscal autonomy, transforming intelligence operations into a self-funding revenue stream for the regime.
- **Edge-device exploitation priority:** FoolishSlug led North Korea’s effort to exploit unpatched edge devices. In early 2025, they were among the first North Korean actors to weaponize critical vulnerabilities in VPN concentrators and secure gateways, using them as initial entry points to move laterally into isolated operational technology (OT) networks.
- **LotL mastery:** To avoid detection in sensitive networks, they almost exclusively used legitimate administrative tools like PowerShell, PsExec, and NetScan. Their 2025 playbook focused on making their lateral movement indistinguishable from a standard network audit, enabling them to remain undetected in victim networks for extended periods of time.



Iran: High-level credential theft and lateral movement via compromised government infrastructure

In general, Iran tightly integrates digital espionage with kinetic military objectives. By leveraging a human-in-the-loop social engineering model and weaponizing the inherent trust in regional government communications, Iranian actors aim to achieve high-precision intelligence gathering in support of military operations. This was evident in 2025 when groups like MuddyKrill reportedly obtained real-time battle damage assessments via hijacked CCTV streams, while ConvolutedKrill and CrustyKrill exploited cloud-native services and employed ClickFix tactics to bypass modern authentication and maintain persistent visibility into adversary command structures.

We expect this integration of cyber reconnaissance and physical targeting to continue through 2026, as Iranian actors further refine their use of hijacked infrastructure to provide near real-time intelligence for kinetic operations.

Threat actor profile



MuddyKrill



Aliases / AKAs: Muddywater, Mango Sandstorm, Boggy Serpens, Static Kitten



Primary objective

Integration of digital espionage with kinetic military objectives and battle damage assessment



Targeting

Israel, MENA, and European critical infrastructure



Tools and TTPs

BugSleep, LiteInject, StealthCache, Phoenix, and UDPGangster

Strategic profile and context

MuddyKrill is a group that integrates digital reconnaissance with physical targeting, including hijacked CCTV surveillance.

Institutional ties and recruitment

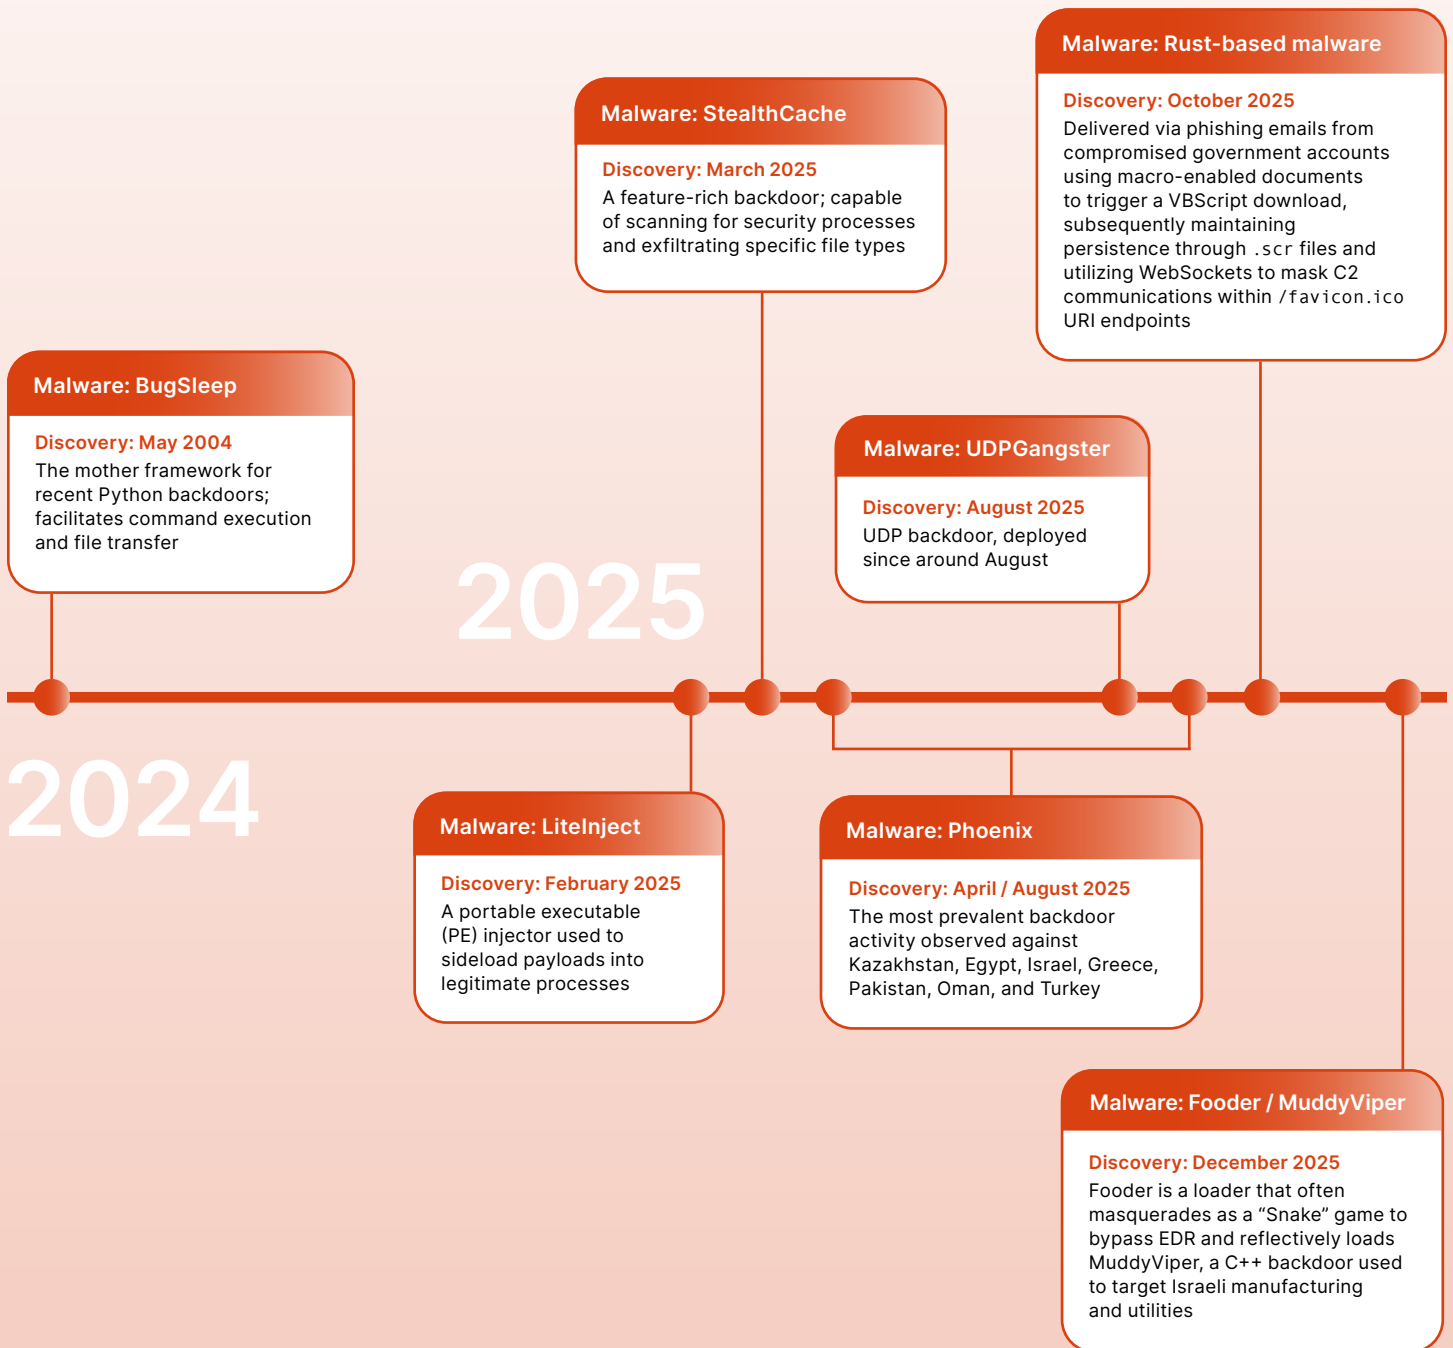
- **Front entities:** Research indicates a strong pipeline between Dadeh Pardazan and the Ravin Academy. In late 2025, a major data leak confirmed that Ravin Academy functions as a primary training ground for cyber operators, specifically feeding talent into MuddyKrill and ConvolutedKrill.
- **Tactical overlap:** Techniques demonstrated in Ravin Academy webinars (e.g., CVE-2020-0688 and CVE-2020-1472 exploits) have appeared in active MuddyKrill campaigns within weeks of the training sessions.

Recent campaigns (2025)

- **Israeli critical infrastructure:** Throughout 2025, MuddyKrill targeted Israeli hospitals, transport, and manufacturing. Notably, in June 2025, the group accessed CCTV streams across Israel to provide real-time battle damage assessment during kinetic missile exchanges.
- **Israeli PM visit to Hungary (Apr 2025):** Capitalizing on the Israeli Prime Minister's (PM) diplomatic visit, MuddyKrill weaponized a Hungarian-language document that contained lures regarding the PM's visit. This themed delivery mechanism was used to deploy SeaSickle malware amid heightened regional tensions.
- **Omani MFA masquerades (Sept 2025):** Targeted global governments in the Middle East and South America (e.g., Columbian Ministry of Foreign Affairs) and South America. This campaign compromised Omani government email accounts and used them to send phishing messages to other government departments in the Middle East and South America. The messages contained malicious attachments that would deliver the Phoenix backdoor.
- **Telecommunications:** Sustained focus on Ethiopian and Moroccan telecoms for the purpose of upstream traffic interception and subscriber monitoring.

MuddyKrill malware arsenal and evolution, 2025

The group has moved away from heavy reliance on PowerShell, favoring custom C/C++, Rust, and Python-based implants to evade detection.



Threat actor profile



ConvolutedKrill



Aliases / AKAs: APT34, Evasive Serpens, Hazel Sandstorm, Helix Kitten, Oilrig, TA452



Primary objective

High-level credential theft and lateral movement via hijacked government infrastructure



Targeting

Middle Eastern government (Iraq MOFA / MOD, Oman) and critical infrastructure



Tools and TTPs

LectureGenie, WebQuestion, ClickFix social engineering, and compromised government email lures

Strategic profile and context

ConvolutedKrill is a credential-theft specialist that leverages compromised government accounts to exploit regional trust between state entities.

Operational focus: Middle Eastern government and critical infrastructure

- **Key TTPs:** High-level credential theft and lateral movement via compromised government infrastructure. The group reuses their C2 HTML pages, with different themes observed.
- **Primary vector:** ConvolutedKrill heavily utilizes compromised government email accounts to send lures to other state entities, leveraging the inherent trust between regional partners.
- **Social engineering lures:**
 - **Fake surveys:** Election or administrative questionnaires hosting .bat files (concealed .NET executables).
 - **Conference lures:** AI-Hol conference invitations used to target regional security and repatriation officials.
 - **ClickFix technique:** Fake Webex or Microsoft Teams error pages that prompt the victim to “fix” the error by executing a malicious PowerShell script.
- **Custom tooling:**
 - **LectureGenie and WebQuestion:** New .NET backdoors utilizing JSON Web Tokens (JWTs) for C2 authentication and sophisticated randomization to mask traffic patterns.
- **Infrastructure:** Extensive use of M247 VPS providers.

Threat actor profile

CrustyKrill





Aliases / AKAs: C5, Smoke Sandstorm, Screening Serpens, TA455



Primary objective

High-touch social engineering and human-in-the-loop interaction for 2FA interception



Targeting

Government and defense contractors via recruitment and career-based lures



Tools and TTPs

PDQ Connect, ClickFix data theft, and fake Google Meet and Microsoft Teams pages

Strategic profile and context

CrustyKrill is a group that uses high-touch social engineering and recruitment lures to compromise victims.

Operational focus: Governments and defense contractors

- **Social engineering (2025):** CrustyKrill employs a human-in-the-loop approach. They utilize fake Google Meet and Microsoft Teams pages, allowing a live operator to interact with the victim in real time to intercept 2FA codes or SMS challenges.
- **Lures:**
 - **ClickFix:** Enables data theft from victim machines and conducts credential harvesting.
 - **Fake careers:** Impersonating Institut Français or Ebix Careers to target professionals.
 - **Remote monitoring and management abuse:** Legitimate remote monitoring and management (RMM) tools like PDQ Connect are deployed under the guise of mandatory software updates.

Cloud exploitation: The group adopts a LotX strategy, hosting C2 on Azure Web Apps (.azurewebsites.net) and using ONLYOFFICE as a payload-hosting platform to blend in with legitimate enterprise traffic.



Cybercrime

From record-breaking DDoS attacks peaking at an unprecedented 31.4 Tbps to the identification of over \$123 million in explicit financial lures, cybercriminals in 2025 demonstrated a relentless move toward high-velocity, automated exploitation. Throughout the entire attack chain, we see how the industrialization of cybercrime contributes to a landscape where individual hackers are replaced by professional-grade “factories.”

This trajectory sets the stage for 2026, where we expect these attack factories to reach even greater levels of industrialization. Whether it is in your inbox, your DMs, or your professional repositories, the goal remains consistent: to steal money or credentials, either as an end in themselves or as the essential means to fuel the broader ransomware and botnet ecosystems.




The phishing factory: From low-friction links to industrialized supply chain sabotage

Technical signals from Cloudflare's Email Security product, which helps block malicious email-borne phishing threats, provide insight into threat actor behavior. We observed that link-based phishing remains the dominant threat vector, favored by both opportunistic criminals and nation-state actors for its extreme economic efficiency.

By leveraging low-friction entry points that bypass traditional EDR and sandbox filters, attackers achieve high efficacy with minimal operational effort. Our screening of 450 million emails reveals a clear hierarchy in the threat landscape: Link-based detections (25%) sit firmly at the top, followed by identity deception (19%) and brand impersonation (16%). To drive victim click rates, attackers lean on these deception tactics to build a false sense of trust, masquerading as a known colleague or a trusted brand to create the psychological green light that leads a user to click. These links are almost always the gateway to credential harvesting (7%), effectively turning a single moment of misplaced trust into a full-scale network breach.

Top threat categories in email detection



Note: Percentages do not add to 100% as emails can have multiple threat categories.

Link-based phishing remains the dominant threat vector, favored by both opportunistic criminals and nation-state actors for its extreme economic efficiency.



This high success rate is driven by a staggering infrastructure gap where a massive volume of mail currently bypasses or fails standard protections.

Our analysis proves that the “open door” [Microsoft warns about](#) is a widespread reality: Nearly 43% of emails failed Security Policy Framework (SPF), over 44% lacked valid DomainKeys Identified Mail (DKIM) signatures, and crucially, 46% failed Domain-based Message Authentication, Reporting, and Conformance (DMARC). This surface area allows phishing-as-a-service (PhaaS) bots to exploit incomplete trust chains, such as the MX-gap vulnerability.

When MX records point to a third-party gateway before reaching Office 365, it creates a blind spot where spoofed external mail may be treated as “internal” or “trusted.” Furthermore, many firms still use DMARC Soft-Fail (p=none or ~all), allowing bots to successfully spoof an organization’s own domain to deliver “urgent” internal notifications.

To maximize these attempts, adversaries wrap malicious links in impersonations of trust-anchor brands like Windows, SANS, and Microsoft. By exploiting the institutional authority of these primary gatekeepers, attackers bypass initial suspicion to target the credentials that serve as the gateway to an organization’s entire digital infrastructure. Other brands that Cloudflare observed attackers frequently

Top 10 most impersonated brands in phishing campaigns

- | | |
|--------------|--------------|
| 1. Windows | 6. Amazon |
| 2. SANS | 7. Instagram |
| 3. Microsoft | 8. Costco |
| 4. Stripe | 9. YouTube |
| 5. Facebook | 10. iCloud |

This list is based on impersonation attempts observed by Cloudflare Email Security.

impersonating include Stripe, Facebook, and Amazon, highlighting a broad mandate across both corporate and personal ecosystems.

Campaign examples: From supply chain sabotage to industrialized PhaaS

Cloudforce One continues to track the shift in how campaigns are executed, moving from simple spam to highly complex supply chain weaponization and industrialized service models. This evolution reflects a broader trend where attackers leverage sophisticated, automated frameworks to achieve maximum impact with minimal manual effort, transitioning from reactive tactics to proactive, large-scale infrastructure abuse.

The Shai-Hulud 2.0 supply chain campaign

The most dangerous evolution of low-effort entry is the Shai-Hulud 2.0 campaign (late 2025), which weaponized the Node package manager (npm) ecosystem through a single phished developer. By gaining access to a maintainer’s npm and GitHub tokens, attackers deployed a self-replicating worm that automatically injects malicious preinstall scripts into every package the victim maintains. This creates a massive ripple effect, potentially infecting thousands of downstream developers and millions of users who rely on these trusted libraries as fundamental building blocks of their own applications.

Beyond simple propagation, the malware is a highly efficient secret harvester. It scans for and exfiltrates credentials for AWS, Azure, and GCP, effectively turning a single laptop compromise into a full-scale cloud infrastructure breach. In a terrifying shift toward extortion through destruction, the 2025 variant contains a dead man’s switch. If the malware detects that its C2 channels are severed or tokens revoked, it triggers a wiper payload that deletes the entire home directory on Linux or %USERPROFILE% on Windows, punishing the organization for attempting to remediate the infection.



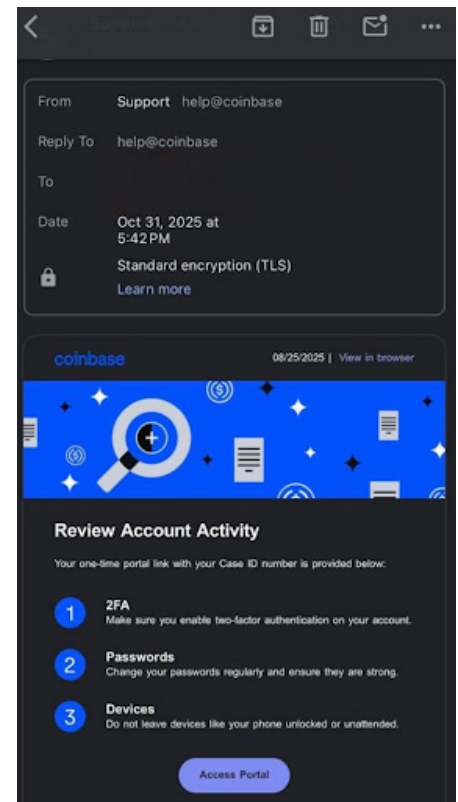
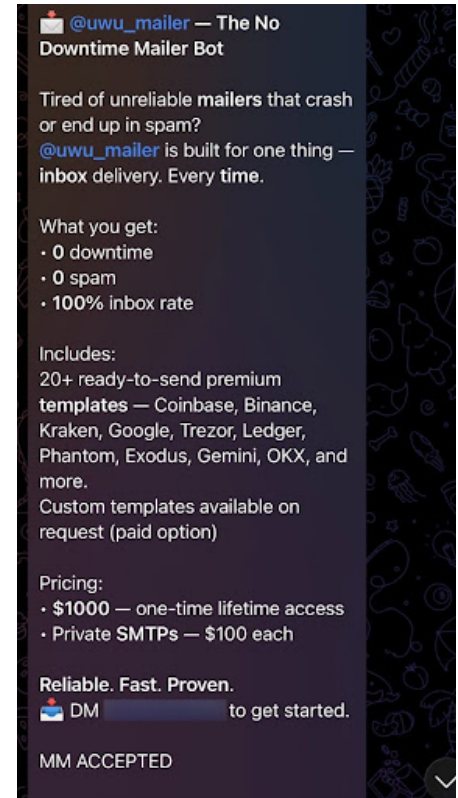
Industrialized PhaaS

Supporting these high-level attacks is the rise of Telegram-based PhaaS modules, exemplified by the recently disrupted [Raccoon0365](#) enterprise. These services go beyond simple spam; they are standardized industrial pipelines marketed as no-downtime mailer bots that significantly lower the barrier to entry for cybercriminals. By offering tiered subscription models — such as \$355 for 30 days — these enterprises provide a plug-and-play marketplace for harvesting credentials, cookies, and data from Microsoft 365, OneDrive, and SharePoint accounts to enable wide-scale financial fraud and extortion.

These industrialized pipelines offer 100% inbox delivery guarantees by using the automated rotation of clean residential proxies and “warm” IP addresses to bypass reputation-based filters and security gateways. They also provide turnkey brand impersonation through high-fidelity, real-time templates for platforms like Google, Microsoft 365, Kraken, and Gemini that are virtually indistinguishable from legitimate logins. To maintain stealth, these kits utilize advanced evasion techniques, including CAPTCHA-based human verification, browser fingerprinting, and anti-analysis scripts that actively disable browser consoles to hide malicious code from security researchers.

Most critically, these bots integrate MFA-bypassing adversary-in-the-middle (AitM) tech. By acting as a transparent proxy between the victim and the legitimate service, these kits harvest live session tokens rather than just static passwords. This effectively neutralizes standard MFA, as the attacker captures the already-authenticated session state.

As seen in Cloudforce One’s Raccoon0365 disruption, these actors often abuse legitimate cloud infrastructure to shield their backend servers, making the attack appear to originate from a trusted network. This shift toward automated, destructive extortion means that traditional perimeter defenses are no longer sufficient; security must now be as fast and adaptable as the automated threats it aims to stop.



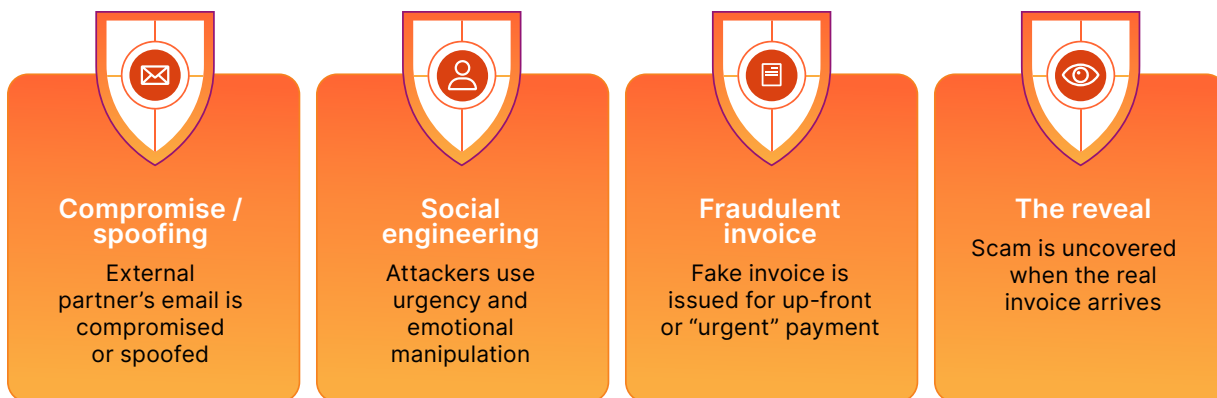
Advertisement for a PhaaS mailer bot on Telegram

The BEC request factory: Intercepting \$123 M in highly targeted financial lures

In 2025 alone, Cloudforce One analysts identified a staggering \$123,455,786 in explicit financial theft attempts. Within this landscape, name impersonation has proven to be the most lucrative tactic for threat actors, accounting for the vast majority of these funds. Our data shows that while the financial requests can vary wildly — with a total range of \$5,567,070 between the largest (\$5,567,520.00) and smallest (\$450.00) attempts — the sweet spot for attackers is remarkably consistent. **The mean financial theft attempt sits at \$49,224.80**, nearly mirroring a median of \$48,950.00 and a mode of \$49,860.00. These figures suggest a calculated strategy where fraudsters aim for amounts large enough to be profitable, yet small enough to potentially bypass more stringent executive approval thresholds.

The true danger of business email compromise (BEC) lies in thread hijacking, where an attacker interjects themselves into an existing, legitimate conversation. This level of intervention is incredibly difficult to detect because it utilizes legitimate communication channels; to automated systems, these requests appear as benign, everyday business activity. **In one standout instance, our team intercepted a single hijacked compromise that saved a client \$5.5 million**, highlighting the critical reliance on manual intervention and expert analysis by teams like PhishGuard to prevent catastrophic losses.

As we move into 2026, we expect attackers to use generative tools to automate thread hijacking at scale, allowing them to maintain this precise ~\$49,000 sweet spot across thousands of concurrent conversations without the need for manual oversight.



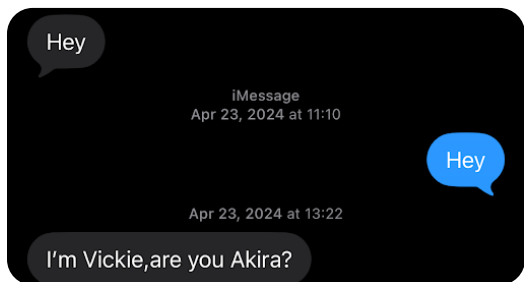
BEC lifecycle

Pig butchering: How scammers weaponize emotional trust and wealth personas to drain life savings

The Cloudforce One team has observed a number of “pig butchering” scams in 2025. Pig butchering (known as shāzhūpán) is a sophisticated long-con investment scam that originated in the People’s Republic of China and is a metaphor for the agricultural practice of fattening up a pig before slaughter; scammers build intense emotional trust with a victim over weeks or months (“fattening”) before stealing their life savings through a fraudulent investment platform (“butchering”).

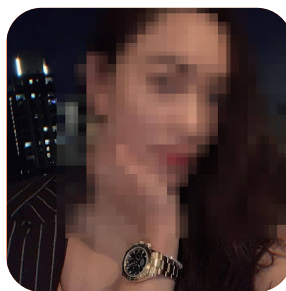
1. A classic “wrong number” introduction

Pig butchering begins with a lure, where scammers initiate contact via “wrong number” text messages or social media to build initial rapport. The scammer often uses a benign “Hey” to bait a response and then apologizes for the “mistake” to launch a conversation.



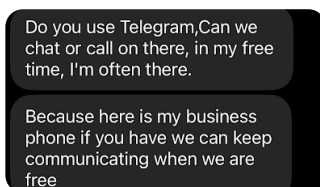
2. The “wealth persona” lure

Following this initial contact, the grooming phase begins. Our analysts have noted the use of wealth personas, where scammers share photos featuring luxury items like high-end watches to establish themselves as successful investors.



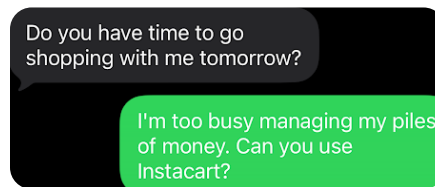
3. Moving to Telegram

To evade security filters and deepen the grooming process, attackers force a platform migration, moving victims to unmonitored messaging apps like Telegram.



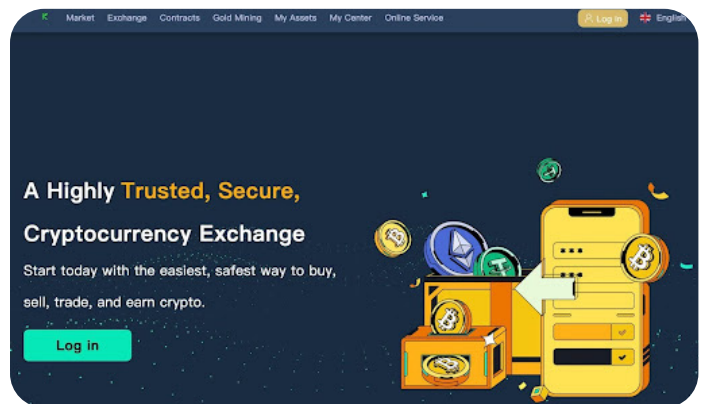
4. Defensive redirection

The example below demonstrates defensive redirection, as the potential victim successfully deflects a grooming attempt. The interaction highlights the strategic ways scammers attempt to integrate themselves into a victim’s daily schedule.



5. A fake exchange

Once trust is solidified, the scam moves to the fraudulent exchange phase. Victims are directed to deposit cryptocurrency into fake exchanges that report deceptive, “amazing” gains to encourage further investment before the victim is ultimately prevented from withdrawing any funds. Victims that request their funds will be asked to send in more funds to cover “fees” or “taxes.”



Recovering any of the stolen money is difficult, if not impossible. In some cases, the cryptocurrency can be traced to an exchange and a legal authority can compel the exchange to return the funds to a victim. In most cases, the funds are quickly laundered so that the legal approach is less effective. Throughout 2026, we anticipate that the grooming phase will become increasingly automated through LLM-powered personas. This artificial intimacy allows a single operative to “fatten” hundreds of victims simultaneously with highly personalized, AI-generated emotional lures that are virtually indistinguishable from human interaction.

The infostealer engine: How stolen logs fuel the global ransomware pipeline

The Cloudforce One team emphasizes that an infostealer infection must be treated with the same urgency as a live network intrusion. These tools have evolved far beyond nuisance malware; they now serve as the primary engine for a transactional criminal supply chain fueled by initial access brokers (IABs). This pipeline begins when IABs purchase logs — bulk collections of stolen data — from infostealer operators, validate the credentials, and auction off high-value corporate access to ransomware cartels.

This evolution has led to staggering corporate exposure; according to Verizon's 2025 Data Breach Investigations Report, 54% of all ransomware attacks in 2025 traced back to infostealer-enabled credential theft.⁵ A critical case in point occurred in early 2025 when the HellCat ransomware group successfully breached major global organizations, including Jaguar Land Rover and Telefónica, specifically by leveraging Jira credentials and session tokens sourced directly from infostealer logs.

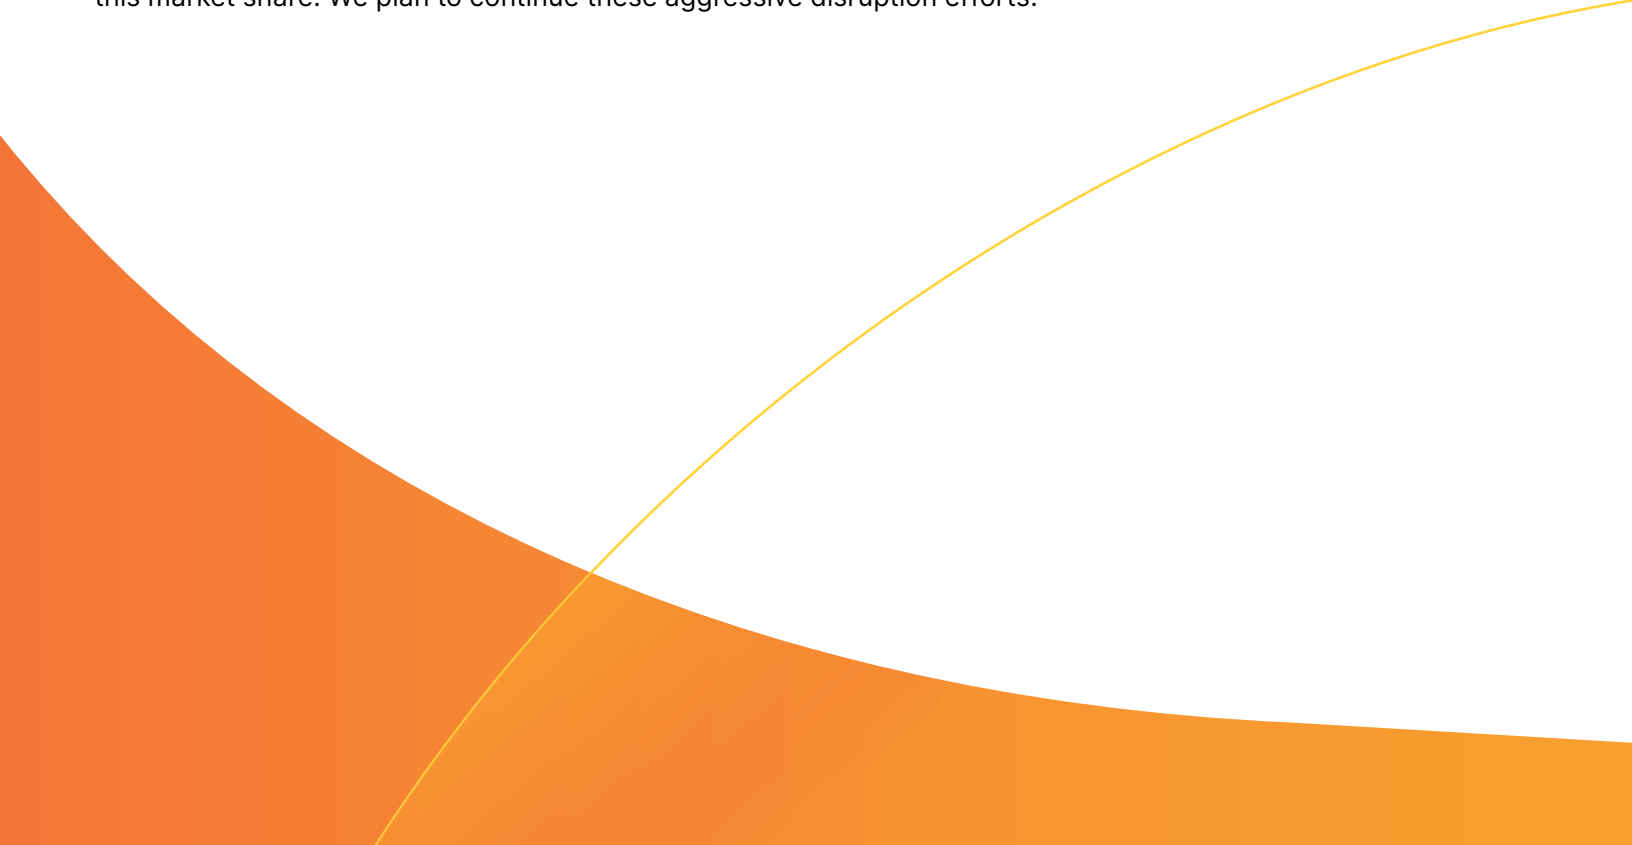
The Lumma disruption and the path ahead

In May 2025, our team participated in a high-impact, coordinated global operation to disrupt [Lumma Stealer](#) (LummaC2), a premier example of the malware-as-a-service (MaaS) model. Lumma represents the industrialization of data theft, providing criminals with customized malware builds and professional-grade dashboards to manage stolen data. Its primary effectiveness stems from the exfiltration of session tokens and active cookies, which neutralizes standard MFA and allows attackers to log in rather than break in.

Cloudforce One's role in this disruption included deploying Turnstile-enabled interstitial warning pages across malicious C2 domains to prevent malware bypass and taking direct action against the accounts used to configure this infrastructure.

While our 2025 focus remained on disrupting these established MaaS giants, the 2026 landscape will be defined by their successors — variants we predict will move toward fully automated token-rotation-bypass to collapse the time between initial infection and full ransomware deployment into a matter of hours.

Cloudforce One is already tracking the emergence of vacuum-fillers and new MaaS variants attempting to reclaim this market share. We plan to continue these aggressive disruption efforts.



Ransomware 2.0: Weaponizing authorized access and human-in-the-loop sabotage

Cloudforce One observes that the modern extortion landscape has shifted from a purely technical encryption challenge into a high-fidelity identity and access crisis. The weaponization of authorized credentials and internal collaborators has become the primary path for high-impact breaches, signaling a move beyond traditional malware toward the exploitation of legitimate access.

This evolution is driven by a sophisticated infostealer-to-extortion pipeline, which in 2026 may transition from an emerging threat to the standard operational baseline for throughput and speed in the ransomware ecosystem. Attack chains frequently begin with malware like LummaC2 harvesting credentials for Citrix, Microsoft RDWeb, and browser-based VPNs. These high-value targets are then handed off to ransomware groups who bypass the difficult intrusion phase entirely, focusing their efforts on data theft and direct extortion. This cycle is being further compressed by AI-accelerated attack cycles; GenAI has shrunk timelines from days to minutes, as threat actors leverage LLMs in real time to rewrite code that bypasses EDR and to map complex network topologies for rapid lateral movement.

The targeting of “critical continuity” has also intensified, with manufacturing and critical infrastructure now representing over 50% of all targeted attacks.⁶ Attackers prioritize these sectors because the immense cost of operational downtime creates a desperate and immediate incentive for high-value payouts. This strategy is often supported by human-in-the-loop access, where attackers recruit remote workers in lower-income regions for as little as \$50 to provide initial access to billion-dollar networks, essentially turning the trusted interior into an active attack vector.

Lastly, in a majority of observed cases, exfiltrated data is used as leverage for ransom while bypassing traditional backup-and-restore defenses that organizations typically rely on to recover from a standard encryption attack. In other words, pure extortion has become the new standard in the ransomware ecosystem.



Case study

The authorized insider investigation and unmasking

A recent Cloudforce One REACT incident response investigation provides a definitive example of ransomware's shift to high-fidelity identity and access crisis. In this instance, a company's trusted employee with high-level permissions leaked sensitive client metadata and source code following a personal grievance and launched a high-value extortion campaign. This quiet crime scene left no traditional malware signatures; instead, investigators had to map a "shadow path" where the insider staged data over several weeks using legitimate production access during standard working hours.

To unmask the threat actor, the team launched a response and investigation rooted in both threat and psychological intelligence indicators. For example, investigators poked the insider in internal communications by downplaying the attacker's skill. This provoked a defensive response in the next ransom email, where the actor used linguistic slips that mirrored the investigators' recent internal comments, narrowing the suspect pool to those with access to specific logs.

By merging technical logs with behavioral science — using sentiment analysis on internal employee communications to match the tone and linguistic patterns of the ransom notes — the team identified a specific individual correlated to threat indicators. Once the company had this forensic evidence in hand, they were able to partner with law enforcement to intercept the employee as they attempted to flee the country. **Accounting for human risk is now just as vital as patching software vulnerabilities.**

The triple-threat bot chain: From identity breach to host compromise to disruption

The sheer scale of automated traffic is a pervasive threat to organizations globally, with approximately [30%](#) of all HTTP traffic Cloudflare observes originating from bots. While some of this is benign — such as search engine crawlers — a significant portion is dedicated to a high-speed attack chain that bridges identity theft, host compromise, and service disruption.

proxy services. This makes malicious traffic look like it is originating from legitimate users, allowing attackers to systematically test defenses from the network layer up to the application layer while remaining below the radar of traditional filters.

Heading into 2026, we expect the use of successor networks like Kimwolf to continue expanding, making this residential-proxy tunneling the default operational baseline for evading IP-based filters.


Bots as the engine of identity exploitation

The chain begins with automated account takeover (ATO), where bots weaponize compromised credentials at industrial scale. This process is fed by a circular ecosystem of harvesting sites and infostealers like Qakbot and Emotet. These bots do not simply look for vulnerabilities in the traditional sense; they exploit human behavior, specifically the fact that [63%](#) of all human logins involve credentials that have already been compromised elsewhere.

Attackers use bots to test stolen username and password pairs across thousands of sites per second. The scale is staggering, as Cloudflare data shows that [94%](#) of all login attempts originate from bots. To bypass modern detection, attackers use tools like Selenium and Puppeteer to mimic human mouse movements and realistic scrolling, allowing them to bypass traditional session intelligence during a credential stuffing assault.

Bot vs. human

Bot (automated) vs. human HTTP requests distribution



Cloudforce One tracks these malicious bots not as isolated scanners, but as a continuous lifecycle. This chain relies on massive, distributed botnets — such as Aisuru and its successor Kimwolf — which mask the source of attacks by tunneling through residential

30% of all HTTP traffic Cloudflare observes originates from bots

63% of all human logins involve credentials that have already been compromised elsewhere

94% of all login attempts originate from bots



Escalating to host compromise and automated harvesting

Once a bot secures a foothold through credential exploitation, the objective shifts to host compromise and the systematic extraction of data. Infrastructure for this stage is provided by botnets like the dismantled 911 S5 or the high-compute Mantis botnet, which utilizes hijacked virtual machines to launch attacks.

This stage of the chain increasingly targets the interfaces of LLMs. These sophisticated bots are designed to interact with LLMs in ways that exploit vulnerabilities in their input handling or output generation, systematically extracting sensitive or proprietary information that was used in the model's training or has been generated by the model itself. This novel form of data theft bypasses traditional network security measures by leveraging the public-facing nature of LLM interfaces, posing a substantial risk to organizations that deploy or rely on these advanced AI systems.

Bot-driven infrastructure disruption and exhaustion


The final link in the chain is the transition to availability threats, where the botnet power used for access is turned into a weapon of destruction. Denial-of-service (DoS) attacks are the most visible expressions of this power, designed to destroy business continuity. In late 2025, the Aisuru botnet reached record-shattering peaks, demonstrating how residential botnets can cripple major digital properties.

By utilizing botnets like Kimwolf, which recently saw over 550 C2 nodes null-routed in early 2026, attackers hide their attacks within legitimate residential traffic to bypass IP-based blocking. Beyond volumetric floods, modern bots target specific high-cost application functions, such as complex search queries, to exhaust a server's CPU and memory and take a site offline with minimal traffic.

This pivot from stealthy exploitation to massive infrastructure bombardment marks a shift in the threat actor's endgame: When access is no longer enough, they weaponize the entire bot chain to achieve total operational blackout. This industrialization of disruption is reflected in the staggering surge of hyper-volumetric activity observed by Cloudforce One.



DDoS attacks by year and type



DDoS attacks observed by Cloudflare

In 2025, the total number of DDoS attacks observed by Cloudflare more than doubled to an astonishing 47.1 million. The most substantial growth was in network-layer DDoS attacks, which more than tripled year over year.

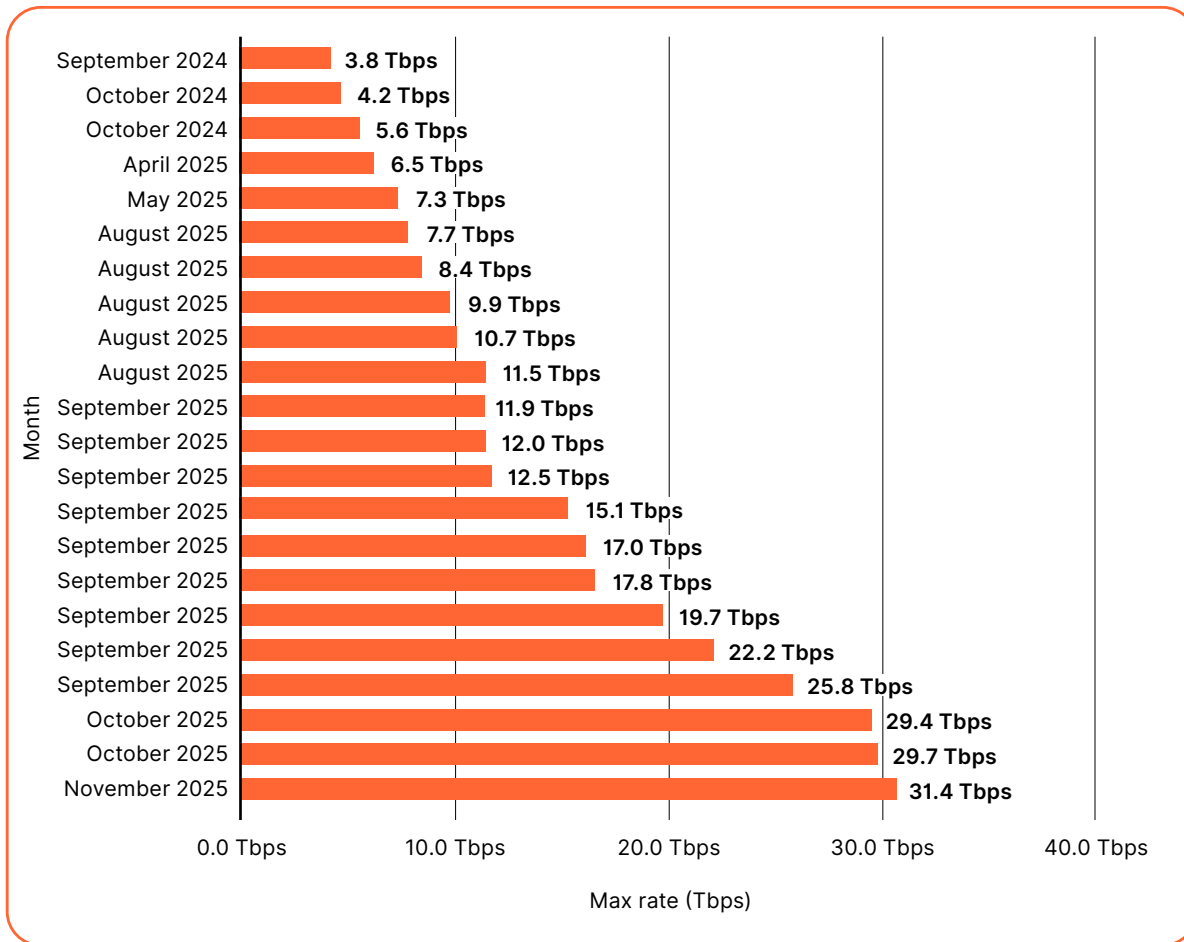
Once viewed mostly as a nuisance, DDoS attacks have become powerful weapons for external actors seeking to disrupt operations and undermine business continuity. On average, Cloudflare mitigated 5,376 DDoS attacks every hour: 3,925 network-layer DDoS attacks and 1,451 HTTP DDoS attacks.

The year was defined by the rise of hyper-volumetric attacks, largely fueled by the emergence of the Aisuru and Kimwolf botnets, which control an estimated 1 to 4 million infected hosts. Hyper-volumetric attacks are now occurring with routine frequency, with Cloudforce One observing 19 new world record attacks in 2025. The current record holder, a 31.4 Tbps attack in November, was a massive UDP flood launched by the Aisuru botnet. This attack was nearly six times the peak volume of 2024's largest attack.

In 2025, the total number of DDoS attacks observed by Cloudflare more than doubled.



World record DDoS attacks



Hyper-volumetric DDoS attacks continue to grow in size and frequency

Heading into 2026, multi-terabit attacks are likely to become the new baseline for targeted campaigns. The democratization of massive botnets like Aisuru means that even mid-tier threat actors can now launch hyper-volumetric attacks that were once the sole province of nation-states.

Organizations must shift from reactive to proactive, autonomous defense postures, as the window for human intervention has effectively closed — most 2025 attacks lasted less than 10 minutes, far too fast for manual mitigation. Furthermore, as geopolitical tensions continue to influence cyber activity, industries like automotive and mining should expect to be caught in the crossfire of state-aligned DDoS campaigns.

For a deeper dive into the current DDoS landscape, explore [Cloudforce One's quarterly DDoS threat reports](#).

IV

Community and regional perspectives


While many threat actors operate globally, regional targeting is often driven by local geopolitical tensions and specific economic assets. The following highlights the primary threat actors and key concerns for the Americas, EMEA (Europe, Middle East, and Africa), and APAC (Asia-Pacific) regions. The actors listed in each region below were chosen for their operational gravity in those areas, based on analyses of where they exert the most pressure, maintain the most specialized infrastructure, or pursue their primary strategic objectives.



Americas

The Americas, particularly the United States, remains the most targeted region globally by volume. Threat activity in the Americas is primarily characterized by highly sophisticated state-sponsored pre-positioning and a pervasive, industrial-scale ransomware ecosystem.


Primary threat actors

CallowDuck 

Alias / AKA:
Scattered Spider

Focus: Identity and access management (IAM) abuse

Context: They are the premier example of the shift from malware to identity abuse. They utilize sophisticated social engineering (vishing) and MFA fatigue to compromise cloud environments. They are highly disruptive to major US enterprise and hospitality sectors (e.g., MGM / Caesars).

SleezyShrew 

Alias / AKA:
APT29

Focus: Government, think tanks, and cloud service providers

Context: This group has shifted heavily toward cloud-based targets, utilizing sophisticated session hijacking and token theft to maintain access to sensitive communications. They focus on long-term persistence (e.g., in Microsoft 365 environments) and supply chain compromises (e.g., SolarWinds).

FrumpyToad 

Alias / AKA:
APT41

Focus: State espionage and private financial gain

Context: Known for software supply chain attacks, they target US healthcare, tech, and retail sectors for both intellectual property and personal profit.

DazedToad 

Alias / AKA:
Volt Typhoon

Focus: Critical infrastructure (energy, water, communications, transportation)

Context: Unlike traditional espionage, this group focuses on pre-positioning for potential future sabotage. They are masters of LotL, using built-in network tools to maintain long-term persistence without triggering traditional malware alerts.

CloyingKrill 

Alias / AKA:
APT33

Focus: Aerospace, defense, and petrochemical sectors

Context: A primary arm of Iranian state-sponsored activity. They target defense contractors to narrow military technology gaps and gather regional intelligence.

Qilin, Akira, and Play 

Ransomware proliferation

Focus: Financial extortion through double extortion (data theft and encryption); leads attacks on the telecom and healthcare sectors

Context: These groups represent the industrialized wing of cybercrime. Akira and Play often exploit unpatched VPNs, while Qilin is known for aggressive name-and-shame tactics.

Key concerns

Identity-first intrusions

Identity abuse has now superseded network exploitation as the primary vector for initial access. Threat actors have moved away from “attacking the box” and are now focused on “attacking the session,” prioritizing credential theft, session hijacking, and the systematic bypassing of MFA. By weaponizing compromised identities, attackers can navigate cloud environments and enterprise systems with the appearance of legitimacy, bypassing the need for traditional malware.

Infrastructure pre-positioning

The shift from data theft to operational persistence in critical utilities suggests attackers are preparing for disruptive events rather than just espionage.

Ransomware service maturity


The professionalization of ransomware-as-a-service (RaaS) means even lower-tier affiliates can execute high-impact attacks using sophisticated playbooks.



Europe, Middle East, and Africa (EMEA)

Europe is the second most targeted region, with 22% of global extortion victims located there. The EMEA landscape remains the primary theater for disruptive cyber operations linked to the war in Ukraine and shifting power dynamics in Eastern Europe and the Middle East.

Primary threat actors

ZapoyShrew 

Alias / AKA:
APT44

Focus: Destructive operations and critical infrastructure sabotage


Context: This group is responsible for grid attacks in Europe and the widespread use of high-profile wiper malware.

MuddyKrill 

Alias / AKA:
MuddyWater

Focus: Regional geopolitics and telecommunications


Context: This group targets government and private sectors across the Middle East (Israel, Saudi Arabia, UAE) and Southern Europe for espionage and regional influence.

IncoherentShrew 

Alias / AKA:
APT28

Focus: Political and military espionage

Context: Heavily involved in disinformation campaigns and leak sites (hack-and-leak) targeting European parliaments and NATO-aligned entities.

RuntZander 

Alias / AKA:
White Lynx

Focus: Information operations (IO) and hack-and-leak campaigns aimed at neighbors of Belarus, including Poland, Germany, Latvia, and Lithuania, with targeting focused on government and media

Context: Closely aligned with Belarusian and Russian interests, this group specializes in stealing and leaking sensitive data to influence local elections or damage political figures.

Key concerns

Targeted critical infrastructure
Organizations in energy, finance, or logistics face an elevated risk of being targeted as part of broader geopolitical pressure campaigns.

Wiper malware spillover
Destructive tools intended for regional conflicts often possess self-propagating capabilities that can inadvertently impact global supply chains.

Political hacktivism
Hacktivists are responsible for almost 80% of recorded incidents in the EU, primarily using DDoS attacks to target public administration and banking.

Industrial extortion
Ransomware accounts for over 38% of incidents in the EU transport sector.

Weaponized disinformation
The combination of data leaks and AI-enhanced disinformation makes it increasingly difficult for organizations to defend their brand reputation during regional crises.



Asia-Pacific (APAC)

The APAC region faces a unique blend of high-tech espionage and large-scale financial theft. Threat activity in APAC is also dominated by actors supporting the Belt and Road Initiative and territorial claims in the South China Sea. There is a heavy focus on maritime intelligence and compromise of core network infrastructure to enable broad-scale monitoring and pre-positioning for future conflicts.

Primary threat actors

DazedToad 

Alias / AKA:
Volt Typhoon

Focus: Critical infrastructure and OT pre-positioning

Context: Utilizes LotL techniques to maintain long-term, undetected access to regional military and utility infrastructure.

WorthlessSlug 

Alias / AKA:
Diamond Sleet

Focus: Cryptocurrency and financial systems

Context: While global, their activity in APAC is intense, targeting regional exchanges and financial hubs (Singapore / Japan) to generate hard currency for the North Korean regime.

SoggyToad 

Alias / AKA:
APT40

Focus: Maritime and naval technology

Context: This group focuses specifically on targets in the APAC region that have strategic relevance to the South China Sea, including regional governments and engineering firms.

ClumsyToad 

Alias / AKA:
Mustang Panda

Focus: Southeast Asian governments and international nonprofits

Context: Focuses on entities involved in South China Sea policy, often using lure documents related to regional summits and diplomatic events.

WimpyToad 

Alias / AKA:
Salt Typhoon

Focus: Telecommunications infrastructure and service providers

Context: Specializes in compromising the core of the Internet — ISP routers and telecom switches — to intercept vast amounts of data at the source.



Key concerns

Network backbone compromise
The targeting of ISPs and network edge devices (Salt Typhoon) means that traffic thought to be secure could be intercepted before it even reaches its destination.

Economic and maritime espionage
High-tech manufacturing and maritime industries are at extreme risk of IP theft aimed at closing competitive gaps in regional industrial capabilities.

Supply chain and IT targeting
The information technology and semiconductor sectors are the most targeted industries in the region.

Stealthy persistence (LotL)
The widespread use of LotL techniques makes it exceptionally difficult to differentiate between a legitimate administrator and a state-sponsored actor.

Summit espionage
Major 2026 diplomatic events (ASEAN in the Philippines, APEC in China) are prime targets for state-backed intelligence gathering.

Drivers of regional overlap

The concept of a regionally isolated threat is becoming more and more scarce. While intent is often regional, the infrastructure and target selection frequently cross continental boundaries. The question of regional overlap in the threat landscape is a core challenge in cyber threat intelligence. While many threat actors are geographically focused for political or cultural reasons, the modern threat landscape shows significant and increasing overlap across the Americas, EMEA, and APAC regions.

We can categorize this overlap into three distinct drivers: geopolitical and state-sponsored campaigns, cybercrime, and the technology supply chain. Below are examples of these regional overlaps and the actor behaviors that drive them.

Geopolitical and state-sponsored overlap

The most sophisticated actors (typically Tier 1 nation-states and threat actors aligned with those states) do not respect regional boundaries because their interests are global.

The big four (China, Russia, Iran, North Korea): These actors maintain distinct “desks” for each region.

- **China-aligned actors (e.g., FrumpyToad, DazedToad):** While heavily focused on APAC (Taiwan / South China Sea), they simultaneously target Americas and EMEA for intellectual property theft and pre-positioning in critical infrastructure.
- **Russia-aligned actors (e.g., IncoherentShrew, SleezyShrew):** Historically focused on EMEA and the Americas, they increasingly overlap into APAC to monitor shifting diplomatic alliances.

Cybercrime big game hunting

Cybercriminal syndicates (RaaS) are almost entirely region-agnostic. They target **revenue, not geography.**

- **Overlap pattern:** A ransomware group might hit a logistics firm in Germany (EMEA) on Monday and an insurance provider in Brazil (Americas) on Tuesday.
- **IABs:** These actors sell access to networks globally. An IAB based in Eastern Europe might sell access to a Japanese manufacturing plant (APAC) to a ransomware affiliate based in the Americas.

Supply chain and shared technology stacks

Because many global enterprises use the same software (Microsoft Office, Salesforce CRM, vSphere, etc.), a single vulnerability creates a global overlap.

The vulnerability sprint: When a zero-day is released (e.g., [React2Shell](#)), actors across all regions rush to scan the global IP space, often with the help of attack surface management tools which are typically helping attackers more than defenders. This creates a noise floor of activity that overlaps every region simultaneously.

Specific examples of regional overlap

Several advanced persistent threat (APT) groups serve as primary examples of adversaries with a global reach, systematically targeting organizations across the Americas, EMEA, and APAC regions. These groups often align their operations with geopolitical objectives or large-scale financial theft.

Key exemplars of multi-regional threat actors include:

<p>NervousToad </p> <p>Alias / AKA: APT27</p> <p>A prominent group that consistently targets headquarters worldwide. Its operations span North and South America, Europe, and the Middle East, focusing on sectors such as aerospace, government, and high tech to conduct long-term espionage.</p>	<p>Lazarus group </p> <p>Alias / AKA: Hidden Cobra</p> <p>A sophisticated group of actors known for their immense geographic reach and diverse motives, including financial gain through cryptocurrency theft and strategic sabotage. It maintains an active presence in North America, Europe, and Asia.</p>	<p>ClumsyToad </p> <p>Alias / AKA: Mustang Panda</p> <p>Highly active group that utilizes diverse techniques like spear-phishing and DLL sideloading across a vast geographic footprint. In 2026, its activity was noted in the United States, Europe, and extensively across APAC countries including Australia, India, Japan, and Vietnam.</p>
<p>SleezyShrew </p> <p>Alias / AKA: APT29</p> <p>A group that operates globally to support Russian foreign policy. It targets government and diplomatic entities across North America, Europe, and other regions opposing Russian interests.</p>	<p>CloyingKrill </p> <p>Alias / AKA: APT33</p> <p>An actor that illustrates a broad operational scope, targeting critical infrastructure in the United States (Americas), Saudi Arabia (Middle East / EMEA), and South Korea (APAC), with a specific focus on the energy and aviation sectors.</p>	<p>LockBit and similar ransomware syndicates </p> <p>While often criminal rather than state-sponsored, these groups function as enterprise cybercriminals with global reach. They target large-scale enterprises across every major economic region, including the Americas, EMEA, and APAC, prioritizing high-value critical infrastructure for double and triple extortion.</p>

The strategic hub phenomenon

Cloudforce One observes certain strategic hubs where regional and global threat interests converge, creating an exceptionally high-density threat environment.

Israel



Global cybersecurity and geopolitical nexus

As a primary theater for both regional conflict and global cybersecurity innovation, Israel faces tri-regional pressure. It is targeted by EMEA-based adversarial neighbors for destabilization, Americas-based stakeholders for geopolitical monitoring, and APAC-based actors seeking to exfiltrate world-class defensive technology and military IP.

Singapore



Financial and logistical nexus

Acting as the pivotal gateway for trade and finance between the APAC and Americas regions, Singapore serves as a global listening post. It attracts the full spectrum of intelligence activity, ranging from state-sponsored financial espionage (seeking insights into global markets) to advanced cybercriminal syndicates targeting its high concentration of liquid capital.

Taiwan



Semiconductor and cross-strait tensions

As the global epicenter for advanced semiconductor manufacturing and a focal point of maritime territorial disputes, Taiwan represents a critical silicon shield. It is under constant pressure from APAC-based actors (China-nexus) focused on technological exfiltration and pre-positioning for potential future blockade scenarios. Simultaneously, it attracts Americas- and EMEA-based stakeholders monitoring regional stability and the integrity of the global hardware supply chain.

Ukraine



Kinetic-digital proving ground

Serving as the frontline for modern hybrid warfare, Ukraine is the world's most active theater for destructive cyber operations. It is relentlessly targeted by EMEA-based state actors (Russia- / Belarus-nexus) for critical infrastructure sabotage and psychological operations. Concurrently, it serves as a primary intelligence collection point for Americas and European allies seeking to analyze the evolution of state-sponsored tactical coordination and new wiper malware variants.

Analytical implications

The regional overlaps detailed in this report have three critical implications for global threat analysis:

- **Cross-regional lateral movement:** A local incident is rarely isolated. Adversaries frequently utilize a breach in an EMEA or APAC branch as a low-friction testing ground or staging point for lateral movement into an Americas-based headquarters. Security teams must adopt a unified global response model to prevent cross-regional escalation.
- **Strategic infrastructure proxying:** Threat actors increasingly utilize geographic proxying to complicate attribution and circumvent geofencing. We frequently observe Americas-based actors leasing EMEA-based VPS infrastructure to target APAC entities. Effective tracking now requires monitoring cross-regional routing patterns rather than relying on simple IP-origin telemetry.
- **Temporal vulnerability exploitation (time-zone arbitrage):** Adversaries strategically leverage regional blind spots by conducting high-intensity activity during the weekends or national holidays of the target region. By initiating operations when local defensive capacity is reduced but the attacker's home region is in a standard business day, they significantly maximize their dwell time and operational success rate.





Recommendations

A roadmap for strategic resilience and an identity-centric enterprise model

The key findings in this report signal that the 2026 threat landscape is defined by the weaponization of identity, the industrialization of SaaS supply chain vulnerabilities, and the emergence of hyper-volumetric, autonomous DDoS strikes that outpace human intervention.

To thrive in this environment, organizations must pivot from reactive, infrastructure-centric defense to a proactive, identity-centric resilience model. The following recommendations provide a high-level roadmap for neutralizing these emerging force multipliers and securing the modern, AI-integrated enterprise.



1

Focus AI security efforts on securing workforce AI usage

Prioritize securing how employees interact with LLMs to prevent AI-assisted navigation by attackers. Implement strict data loss prevention (DLP) for AI prompts and deploy browser-isolated environments for generative AI tools to ensure corporate keys to the kingdom aren't inadvertently leaked into model training sets or captured by infostealers.

2

Transition from MFA to identity-first zero trust

Since infostealers like LummaC2 now harvest session tokens to bypass MFA, organizations must move beyond simple one-time codes. Implement phishing-resistant MFA (FIDO2 / passkeys) and continuous monitoring that invalidates sessions instantly if impossible travel or suspicious device fingerprints (like mouse-jiggling software) are detected.

3

Harden the SaaS-to-SaaS connective tissue

The GRUB1 campaign proves that a single compromise of a trusted integration can create a dangerous ripple effect. Conduct an immediate audit of all SaaS API permissions. Apply the principle of least privilege to integrations, specifically looking for over-privileged read / write tokens in tools like Salesforce, Slack, and GitHub that could allow an attacker to pivot between clouds.



4

Implement human-in-the-loop verification for remote hiring

To counter the industrialized North Korean insider threat, move away from purely digital onboarding. Use zero trust biometric verification for all remote video interviews and enforce strict hardware-based geofencing. Corporate laptops should be cryptographically paired to the user's physical location to neutralize "laptop farm" facilitators.

5

Adopt autonomous, hyper-volumetric DDoS defenses

With the Aisuru botnet pushing attacks to a 31.4 Tbps new baseline, the window for human intervention has closed. Organizations must shift to automated, edge-based mitigation that can respond in seconds. Legacy scrubbing center models are no longer sufficient for attacks that peak and conclude within 10 minutes.

6

Isolate peripheral infrastructure to contain exposure

To establish a robust defensive posture, organizations must implement a strategic shift in how they manage IaaS and SaaS dependencies. Specifically, subsidiary and supporting services should operate independently, utilizing dedicated domain names, unique IP addresses, and, where feasible, distinct autonomous system numbers (ASNs).

Eliminate email blind spots with AI-first security

PhaaS bots can rapidly bombard organizations with emails leveraging polymorphic tactics that bypass legacy secure email gateways. Organizations must adopt AI-first email security capable of interpreting these shifting variables and adapting to both incoming and lateral threats. By utilizing signals beyond the email inbox, these systems can better identify and neutralize internal compromised accounts in real time.

7

Cloudforce One overview



Threat intelligence

Gain access to Cloudflare's global threat visibility with a powerful package of threat intelligence tooling and expertise to make your organization smarter, more responsive, and more secure.

[Learn more](#)



Managed defense

Receive 24 / 7 expert monitoring and proactive threat detection across your Cloudflare stack, ensuring continuous protection and rapid response to issues.

[Learn more](#)



Cyber response and readiness

Deploy hands-on experts to neutralize active threats, conduct forensic analysis, and minimize operational impact.

[Learn more](#)



[Contact](#) **Cloudforce One security experts**


[Subscribe](#) to our threat research


Modernize your security with Cloudflare


One security platform. Network to cloud. Apps to AI.



Cloudflare’s one security platform scales seamlessly from network to cloud, and protects all users and data across applications and AI. It leverages an intelligent and programmable global cloud network to unify a range of diverse technologies, harnesses AI-powered threat intelligence to deliver low-touch, high-efficacy protection, and AI-powered data loss prevention detections to distributed organizations of all sizes.

 Secure inbound and outbound traffic from network to application layer, to increase consistency and coverage across distributed IT environments with existing teams and budget.

 Secure AI workloads from models to inference, to increase AI adoption with secure, reliable inference and protected, clean data for trusted AI models.

 AI-powered protection and observability fueled by our global network scale reduces the mean time to detect, respond, and mitigate future critical incidents.

Cloudflare offers an integrated approach to security modernization, simplifying deployment, enhancing visibility, and reducing the TCO while eliminating the multiple vendors and tools required with traditional, fragmented solutions. By delivering security close to the end user and requested resource, Cloudflare ensures resiliency and broad protection against evolving threats.

The result: One security platform. Network to cloud. Apps to AI.

With Cloudflare, customers can realize:

24%
breach risk reduction across the enterprise

35%
time savings on security and DNS management

250
hours of avoided downtime

227%
ROI of Cloudflare over three years

Source: Forester TEI study

Endnotes

1. <https://www.nhpr.org/2025-11-21/russian-hacking-suspect-wanted-by-the-fbi-arrested-on-thai-resort-island>
<https://therecord.media/russian-arrested-thailand-allegedly-void-blizzard-apt-member>
2. <https://hackread.com/us-telecom-breaches-firms-chinese-salt-typhoon-hackers/>
<https://www.reuters.com/world/asia-pacific/china-hacked-email-systems-us-congressional-committee-staff-ft-reports-2026-01-08/>
<https://www.microsoft.com/en-us/security/blog/2025/07/22/disrupting-active-exploitation-of-on-premises-sharepoint-vulnerabilities/>
3. <https://my.f5.com/manage/s/article/K000154696>
<https://cloud.google.com/blog/topics/threat-intelligence/brickstorm-espionage-campaign>
4. <https://dev.to/lovestaco/go-undercover-code-obfuscation-with-garble-1ld6>
5. <https://www.verizon.com/business/resources/reports/dbir/>
6. <https://industrialcyber.co/reports/half-of-2025-ransomware-attacks-hit-critical-sectors-as-manufacturing-healthcare-and-energy-top-global-targets/>



Get in touch with us to learn how to modernize your security.



1 888 99 FLARE



enterprise@cloudflare.com



www.cloudflare.com

This document is for informational purposes only and is the property of Cloudflare. This document does not create any commitments or assurances from Cloudflare or its affiliates to you. You are responsible for making your own independent assessment of the information in this document. The information in this document is subject to change and does not purport to be all inclusive or to contain all the information that you may need. The responsibilities and liabilities of Cloudflare to its customers are controlled by separate agreements, and this document is not part of, nor does it modify, any agreement between Cloudflare and its customers. Cloudflare services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied.

© 2026 Cloudflare, Inc. All rights reserved. CLOUDFLARE® and the Cloudflare logo are trademarks of Cloudflare. All other company and product names and logos may be trademarks of the respective companies with which they are associated.