LUMEN®

BLACK LOTUS LABS® by Lumen

The 2026 Lumen Defender Threatscape Report

Why visibility at breach misses the plot

Table of contents

Executive summary 03

What security leaders need to know 06

Threats deconstructed 16

Kimwolf: The distributed-denial-of-service botnet that rose out of Aisuru 17

Rhadamanthys: A malware-as-a-service platform built like a startup and run like a crime syndicate 23

Brute force attacks: A classic heist trick, rebuilt for scale 27

SystemBC: The high-bandwidth proxy crew built for volume 30

DanaBot: The malware delivery crew that ran like a franchise 33

5socks botnet: The long-running proxy heist hiding in plain sight 37

J-magic campaign: The invisible backdoor living on routers 40

Secret Blizzard: How a second crew hijacked the control room and rewrote the job mid-heist 43

NSOCKS botnet: The proxy crew that turns home routers into disguises 47

Raptor Train botnet: The nation-state botnet with an enterprise-grade command center 50

2026 threat predictions 55

Prediction 1: Targeting will focus more on opportunity than vertical 56

Prediction 2: Setup gets faster as adoption of generative AI and agents goes mainstream 57

Prediction 3: The real signals lie in the network 58

Prediction 4: The best disguises will be legitimate infrastructure 58

Defense guidance: Stopping the heist before it happens 61

1. Defend the edge like it's the vault door 62

2. Shift from indicators to infrastructure awareness 63

3. Treat proxy networks as active threat infrastructure 63

4. Assume blurred lines between crime and espionage 64

5. Use scale against the attacker 64

Conclusion 65

Inside Lumen global internet backbone visibility 66

Research methodology 68

How is a risk determined? 69

LUMEN®

Executive summary

Think of modern cybercriminals and nation-state actors as their own version of an elite heist crew. Instead of cracking safes, they're assembling malicious proxy networks and accelerating operations with generative AI. For business leaders, especially CISOs and cybersecurity leaders, it's now critical to gain upstream visibility into cyber operations before attackers reach the perimeter.

Modern cyber operations look less like isolated break-ins and more like carefully staged heists. Long before a breach makes headlines, threat actors invest significant time assembling the infrastructure that will power their campaigns—using generative AI to continuously scan for exposed edge devices, validate stolen credentials, stand up proxy networks, and test command-and-control (C2) paths. Most organizations only see these operations once the attack reaches their perimeter or endpoint. By then, the preparation is already complete.

Lumen sees these campaigns much earlier.

Backed by our threat research and operations arm, Black Lotus Labs®, Lumen operates from its own vantage point—inside a global internet backbone. Rather than relying on post-infection signals from endpoints or perimeter devices, we analyze backbone-level telemetry to identify coordinated infrastructure behavior as it emerges.

What are edge devices?

Edge devices, such as routers, switches, firewalls, VPN gateways, refer to internet-facing infrastructure and services that sit at the boundary between an organization's internal environment and the public internet, providing access, routing, or control, but typically operating outside traditional endpoint security visibility.

Who is Black Lotus Labs?

Black Lotus Labs is the threat research and operations arm of Lumen, combining unmatched network visibility with expert research, machine learning, and automation to conduct original threat discovery and uncover threat actor infrastructure. We started Black Lotus Labs with the mission of leveraging Lumen's network visibility to help protect customers and keep the internet clean.

When malicious infrastructure is confirmed, Black Lotus Labs disrupts it by blocking, taking down, and notifying on confirmed C2 servers. We then feed this intelligence into Lumen Defender security solutions to block threats before they reach your organization.

Learn more about Black Lotus Labs and Lumen global internet visibility →

LUMEN®

This includes early-stage signals such as large-scale scanning, credential-validation activity, botnet enrollment, proxy formation, and rapid C2 rotation—often days or weeks before those same Internet Protocol (IP) addresses or domains are observed targeting any single enterprise.

From this comprehensive viewpoint, Lumen has identified a major shift for 2026: **modern cyberattacks are increasingly driven by exposure, with threat actors optimizing their targeting for vulnerable edge devices and services.**

Savvy threat actors recognize that it's harder to breach systems that are protected by well-developed endpoint detection and response (EDR) products. Instead, they're maximizing their chances by targeting any internet-exposed edge device or service. Endpoint detection is still a crucial source of intelligence. However, organizations must also consider exposure at the edge—going beyond firewalls into complex end-to-end network visibility. That's why Lumen has decided to launch an annual threatscape report to provide a unique viewpoint into the shifting nature of attacks.

As one of the world's largest internet backbone operators and number one most deeply peered network, our transit traffic provides visibility into 99% of all public IPv4 addresses.[1] We monitor more than 200 billion NetFlow sessions and DNS queries, 2.3 million unique threats, and 46,000 C2s daily. We then use this visibility to spot and block the threats that others can't see—executing 5,000 C2 disruptions in 2025 through takedowns and notifications to proactively protect our customers and secure network infrastructure.

Inside Lumen global internet visibility

- Visibility into 99% of public IPv4 addresses
- Daily monitoring of 200B+ NetFlow sessions and DNS queries and 46,000 C2s
- Daily tracking of 2.3M unique threats
- 5000+ C2s disrupted in 2025

Lumen upstream visibility fundamentally changes how threats are detected and disrupted. New adversary infrastructure has no reputation by definition, but it does exhibit behavior. By correlating how infrastructure communicates, how quickly it is created and abandoned, and how it adapts when portions are blocked, Black Lotus Labs links newly stood-up assets back to emerging campaigns while they are still forming. This campaign-level perspective allows defenders to move beyond chasing individual indicators and instead focus on the systems attackers rely on to operate at scale.

LUMEN®

Our 2026 threatscape report reflects what this vantage reveals. Across criminal and nation-state activity, Lumen has observed a decisive shift toward infrastructure-centric operations: proxy networks becoming a pillar for all manner of attackers, edge devices becoming preferred footholds, and generative AI accelerating how quickly attackers assemble and regenerate their tooling. These trends compress defender response windows and render purely reactive controls insufficient.

This report is not a catalog of breaches. It is a view into how modern attacks are built, where their earliest signals emerge, and why defenders who can see and act at the infrastructure layer hold a decisive advantage in 2026 and beyond.

-Natnael Habtesion, Senior Vice President and Chief Security Officer, Lumen

LUMEN®

What security leaders need to know

How to spot the attack before it starts

Modern cyber operations don't begin at the endpoint or at perimeter security solutions like firewalls or web application firewalls (WAFs). They begin upstream, as threat actors scan for unpatched vulnerabilities and end-of-life (EoL)/end-of-support (EoS) devices and assemble the infrastructure for their next campaign. Here are the top 2025 trends that security leaders need to know and what matters most in 2026:

Top 2025 trends:

Generative AI has changed the tempo: Threat actors iterate and regenerate infrastructure at AI and machine speed. Without upstream intelligence, existing defense systems fall short because they rely on static indicator of compromise (IOC) lists, which are updated infrequently. Today's defenders need intelligence that can identify attacker infrastructure in formation, not human-speed defense systems that only spot attacks after they've been triggered in the enterprise.

Attackers are shifting from endpoints to the edge: Endpoint detection and response (EDR) solutions were widespread in 2025, with 91% of organizations deploying an EDR that covered 72% of in-scope devices on average.[2] This well-established defense pushed attackers to shift to less visible infrastructure—including edge devices with limited forensic capabilities. Multiple government agencies have noted this growing trend, which underscores the importance of broader behavior-based detection. behavior-based detection.[3]

Infrastructure is the new early-warning layer: IP movement, proxy formation, and C2 rotations are the real telltale signs that a new campaign is forming, as seen in the shift from Aisuru to Kimwolf.

Malware-backed proxy networks break traditional trust signals: Hijacked home routers and Virtual Private Server (VPS) proxies let attackers impersonate legitimate users and bypass geolocation and Zero Trust controls.

Attribution is blurring: Secret Blizzard and NSOCKS highlight how nation-state and criminal groups now share, steal, or rent the same infrastructure—meaning context cues like infrastructure relationships and lateral movement matter more than labels.

LUMEN®

Total malicious scanners detected per day (2025)

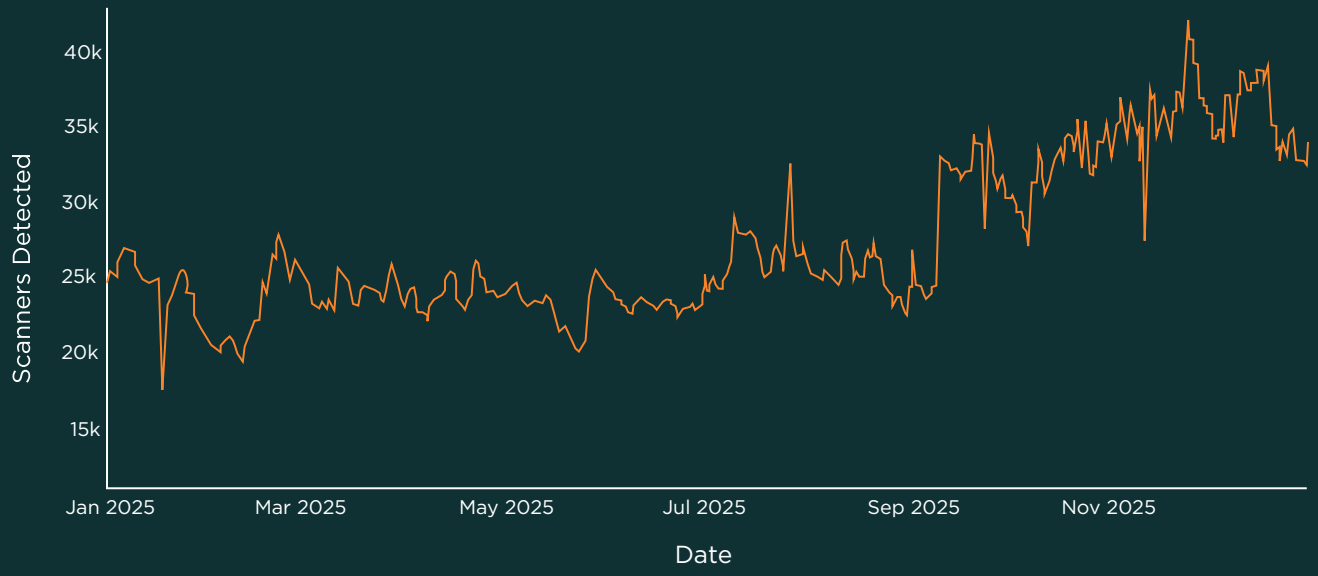

The massive volume of malicious scanning indicates attackers' commitment to taking a systematic, continuous inventory of internet-facing devices and services. Botnet operators, initial access brokers, and advanced actors use this scanning to assess the opportunity space for their exploits on edge devices before they ever launch an attack.

Top enterprise edge devices brute forcing events tracked by Black Lotus Labs (October–December 2025)

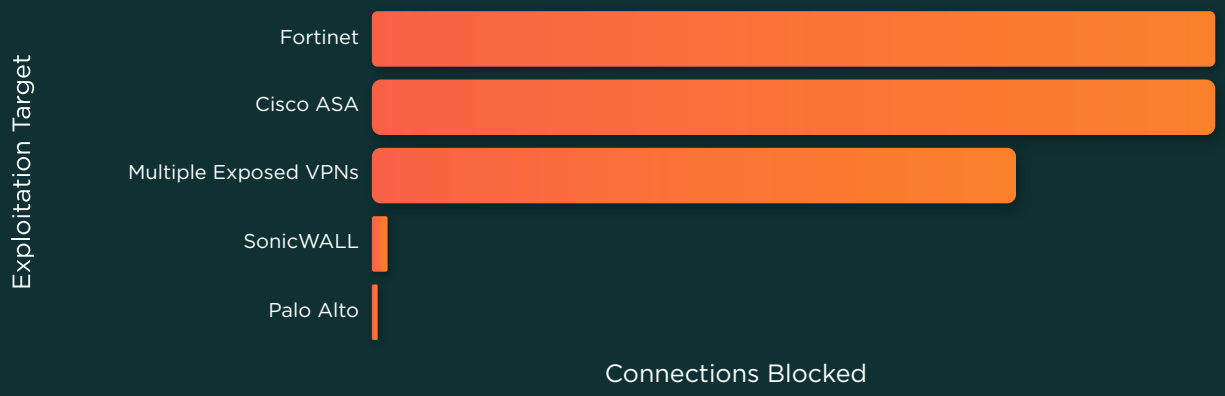

Using Black Lotus Labs' brute force detections, we blocked significant traffic from attack attempts in Lumen Defender Essentials and Plus against enterprise edge devices. This chart represents the five most-targeted devices based on connections blocked by Lumen Defender Essentials & Plus between October–December 2025.

LUMEN®

Dedicated enterprise edge exploitation servers detected per day

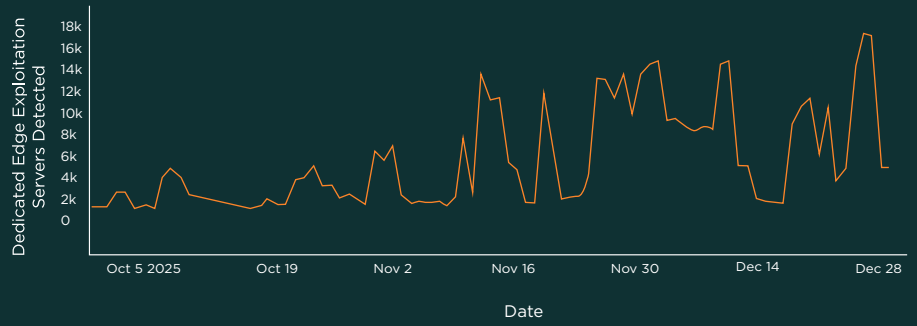

There are several variations of brute force attacks. Thanks to generative AI and distributed computing, adversaries are leveraging bots to generate massive brute force volume. Black Lotus Labs detects these attacks by monitoring specific edge devices and services across our global NetFlow to identify early campaign indicators and automatically block new attack waves for Lumen Defender customers.

2026 threatscape predictions

Attack setup will accelerate to machine speed: Generative and agentic AI will automate scanning, exploitation, and lateral movement—compressing the time between exposure and impact.

Exposure will define risk: Continuous scanning of internet-exposed and end-of-life devices will drive opportunistic, AI-accelerated attacks across all sectors.

The most important signals will live in the network: As attackers favor edge devices with limited forensic visibility, detecting coordinated infrastructure behavior will be essential.

Legitimate infrastructure will become the best disguise: Proxy networks, Small Office/Home Office (SOHO) botnets, and hijacked VPS environments will allow attackers to blend seamlessly into normal internet traffic.

Bottom line? Edge-focused attackers seek to abuse identity management controls by acquiring valid credentials and leveraging them to breach internet-exposed enterprise devices and services. Threat actors also capitalize on mis-configured edge devices and rapidly leverage proof of concept (PoC) exploits for new edge device vulnerabilities before defenders have time to patch.

To stay ahead of these growing threats, security leaders must expand intelligence beyond current investments in perimeter and endpoint security, vulnerability and attack surface management tools, Zero Trust, and SOCs by adopting network infrastructure intelligence and capabilities to proactively block attacks upstream.

Additionally, security teams must reinforce cybersecurity fundamentals, which are still critical in the age of AI. This includes practicing Zero Trust, defense in depth principles like comprehensive logging and monitoring, strict privileged access control, monitoring for privilege escalation, timely patching of exposed systems, and accurate inventories that enforce the retirement of unsupported and EoL devices before they become entry points for adversaries.

LUMEN®

2025: Year in review

Cyber threats in 2025 were shaped less by individual malware families and more by the infrastructure attackers built, borrowed, and hid within. Obfuscation networks and botnets became foundational tools, with multi-layered chains of compromised routers, SOHO devices, Internet of Things (IoT) hardware, and virtual servers blending malicious activity into everyday traffic. Their constant churn complicated attribution and overwhelmed traditional detection models.

From inside the global internet backbone, Black Lotus Labs observed a clear escalation in both scale and intent. Criminal and nation-state actors alike relied on millions of vulnerable, often EoL devices to create moving targets. SOHO and IoT devices allowed attackers to "live off the land" while disguising operations behind residential IP space. Shared use of compromised infrastructure further blurred the lines between criminal and state activity, increasing the risk of missed detections and delayed response while allowing hackers to evade prosecution.

This shift was especially visible in how attackers targeted access. Throughout 2025, Black Lotus Labs observed persistent focus on critical infrastructure, government, military, telecom, managed service providers, IT vendors, communications systems, and academia. Initial access increasingly relied on stolen credentials used through legitimate pathways like VPNs, followed by minimal on-host persistence and aggressive log deletion. At the same time, attacks against edge devices, management systems, and hypervisors surged—high-value assets with broad administrative reach but limited forensic depth.

As seen in disruptions like DanaBot during Operation Endgame II, the 5socks botnet, and the Lumma infostealer, the global internet backbone has become both the staging ground and the battleground. From this vantage point, adversaries operate like cloud-native operations: botnets function as platforms, C2 frameworks mimic enterprise applications, and proxy markets act as global exchanges for access and anonymity. Whether it was a home router or an enterprise firewall, if a device was visible and vulnerable, attackers had the time and incentive to compromise it.

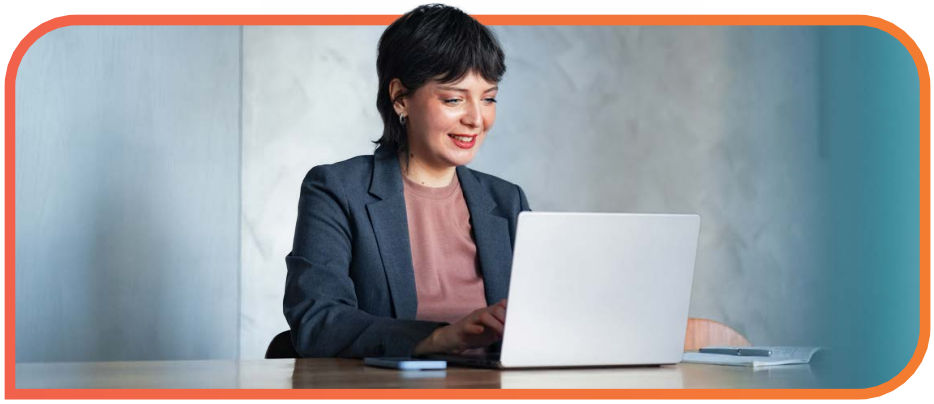

The following six insights capture how these operations took shape throughout 2025.

1 Generative AI changed the tempo, unleashing attacks at machine speed.

The cyber criminals and nation-state groups of 2025 didn't rely on human operators pulling levers manually. Their infrastructure evolved independently.

Threat actors embraced automated tasking, scanning, infrastructure rotation, and C2 management, in addition to adopting new generative AI-enabled attack vectors. Botnets like Raptor Train used queued job systems that assigned tasks like a warehouse logistics engine. Proxy networks like SystemBC relied on auto-relay infrastructures that rerouted operations with no human involvement. This automation sustains malicious campaigns indefinitely with minimal oversight. Human operators step in only to update exploits, payloads, or targets, letting machines handle daily maintenance. Even these human-operated tasks are increasingly assisted by generative AI tooling.

Threat actors are also getting more creative in how they leverage different types of devices. For example, SystemBC is adept at using high-bandwidth VPSs and shifting them into additional operations. This scalability and adaptability pushes organizations to rethink how they detect and disrupt attacks, examining new types of signals to anticipate threat actors' movements before they breach the network perimeter and pre-positioning IP address blocks to stop threat actors in their tracks.

2 Attackers are moving deeper into the network itself—hiding in the infrastructure.

In 2025, we saw sophisticated adversaries shift more of their attention away from endpoints to vulnerable, less-secured devices at the edge, like routers, VPN gateways, and firewalls. Here, they can "live in the middle," stealing credentials to access the protected information beyond, hiding their activity inside the connective tissue of the internet. Sophisticated attackers know that completing the heist in a single sitting may alert defenses. By lurking in devices outside of the reach of standard security controls and reaching back out days, weeks, or even months after the initial access, attackers can better evade detection and prevent defenders from connecting the dots.

For example, J-magic backdoor attacks planted a passive listener onto enterprise-grade Juniper routers—likely with the intention of gaining discreet access, intercepting credentials, and positioning passive listening malware within corporate networks. Once embedded, J-magic installs an agent to passively scan for predefined parameters before activating and issuing a secondary challenge. J-magic shows how savvy threat actors can embed deep within networks and establish reverse shells on local file systems—enabling operators to control devices, steal data, or deploy malicious software.

3 Criminal ecosystems professionalized, adopting the polish of legitimate SaaS.

Gone are the days of crude panels and chaotic infrastructure. In 2025, cybercriminal operations became indistinguishable from professional software businesses.

Rhadamanthys, for example, rebranded itself as RHAD Security, complete with customer support; subscription tiers; a dedicated, curated sales platform; and a reverse-proxy network add-on. Its operators infected hundreds of thousands of victims worldwide and ran hundreds of daily C2 servers. More than 60% of Rhadamanthys C2s remained undetected on VirusTotal at the time of Black Lotus Labs's initial publication.

4 Malware-backed proxy networks became full-fledged economies of disguise.

Adversaries can spend months, sometimes even years, laying the groundwork for their next operation. That's why we saw threat actors industrialize their entry and escape routes in 2025.

Compromised home routers, IoT devices, and high-volume VPS hosts became rentable identities—on-demand personas that can be purchased for a few dollars' worth of cryptocurrency. For example, the 5socks botnet transformed IOT and SOHO devices in the residential IP space into a sprawling anonymity market—capitalizing on the massive pool of EoL consumer devices that are easily re-exploited and repurposed into criminal proxy networks. This tactic allowed attackers to capitalize on the surge in remote work following the Covid-19 pandemic and disguise themselves as trusted remote employees, bypassing geofencing, autonomous system numbers (ASN)-based blocking, IP reputation checks, and many Zero Trust location signals.

With 12% of U.S. workers fully remote in 2026 and 27% hybrid, this trend underscores the importance of upstream visibility for identifying compromised devices before they impersonate corporate users.[4]

5 Nation-state & criminal crews blurred together on shared infrastructure.

In classic heist films, every crew has its signature—the safe expert, the hacker, the criminal mastermind. But in 2025, those signatures overlapped. Nation-states piggybacked on criminal infrastructure; cybercriminals reused tooling forged by intelligence services.

Secret Blizzard, a Russian-based threat actor also known as Turla, infiltrated infrastructure built by other nation-state actors, moving stealthily through Iranian and Pakistani C2s to deploy malware and acquire data. Raptor Train (state-linked), Secret Blizzard (state-linked), and 5socks (a criminal botnet) often operated in the same digital neighborhoods, using the same devices, ports, and sometimes even hosts to disguise their movements and make it difficult for defenders to link activity to specific threat actors.

Attribution became less about ownership and more about intent. Everyone was using the same or similar digital playgrounds to achieve their distinct goals while making it difficult for defenders to know exactly who they're protecting against.

6 The global backbone transformed from a conduit into an early-warning system.

2025 underscored how threat actors are constantly exploring new attack vectors and using AI to accelerate the speed and sophistication of their operations. This transforms the global internet backbone from an invisible highway for threat actor transport into a critical detection and disruption layer.

As adversaries weaponize routers, appliances, and management planes, the infrastructure layer itself becomes the first theater of operation. Backbone telemetry exposes patterns that endpoint logs alone can't catch, from the quiet appearance of a new proxy node to botnet recruitment, automated scanning, or the sudden regeneration of C2 servers as operators prune and replace infrastructure at machine speed.

These signals emerge upstream in the connective tissue of the internet, turning backbone operators like Lumen into frontline defenders. Through Black Lotus Labs, we null route infrastructure to degrade adversary capability and enable coordinated cross-industry actions. Lumen Security customers don't have to wait for global disruption on select campaigns to realize Black Lotus Labs defenses for their networks, as they received minute-to-minute updates of Black Lotus Labs threat intelligence from the global IP backbone integrated into their solutions and services.

Together, these insights surface a critical shift for defenders in 2026: the decisive battleground is no longer just the endpoint—it's the infrastructure beneath.

LUMEN®

2025's most wanted

Black Lotus Labs detected over 2 million IPs associated with very large botnet families in 2025 such as Aisuru and Vo1d. On any given day, the largest Black Lotus threat families detected by Black Lotus Labs in Lumen's global telemetry were Aisuru-based proxies and bullet proof anonymization platforms.

Top 10 threat families based on indicators detected (2025)

	Family	IPs Detected
1	Aisuru	2,948,616
2	Vo1d	2,519,125
3	AWM	2,356,202
4	Aisuru Proxies	2,246,215
5	NexusNet	2,231,627
6	NSOCKS	1,482,113
7	StarVPN Proxies	1,428,920
8	Mylobot	1,377,437
9	SOCKS5 Systemz	1,184,444
10	Tofsee	1,132,114

Top 10 threat families based on average daily bot count (2025)

	Family	Average IP Count
1	Aisuru Proxies	129,487
2	Mysterium	45,097
3	Aisuru	31,549
4	Vo1d	28,906
5	AWM	18,310
6	NexusNet	16,399
7	Mylobot	12,035
8	NSOCKS	11,881
9	RemProxy	11,451
10	StarVPN Proxies	11,035

The IPs detected for the top 10 threat families include bots and infrastructure such as C2s, while the second chart shows the families that maintained the largest pool of bots in Lumen visibility in 2025.

LUMEN®

Top 10 threats blocked in Lumen Defender based on Black Lotus Labs tracking (2025)

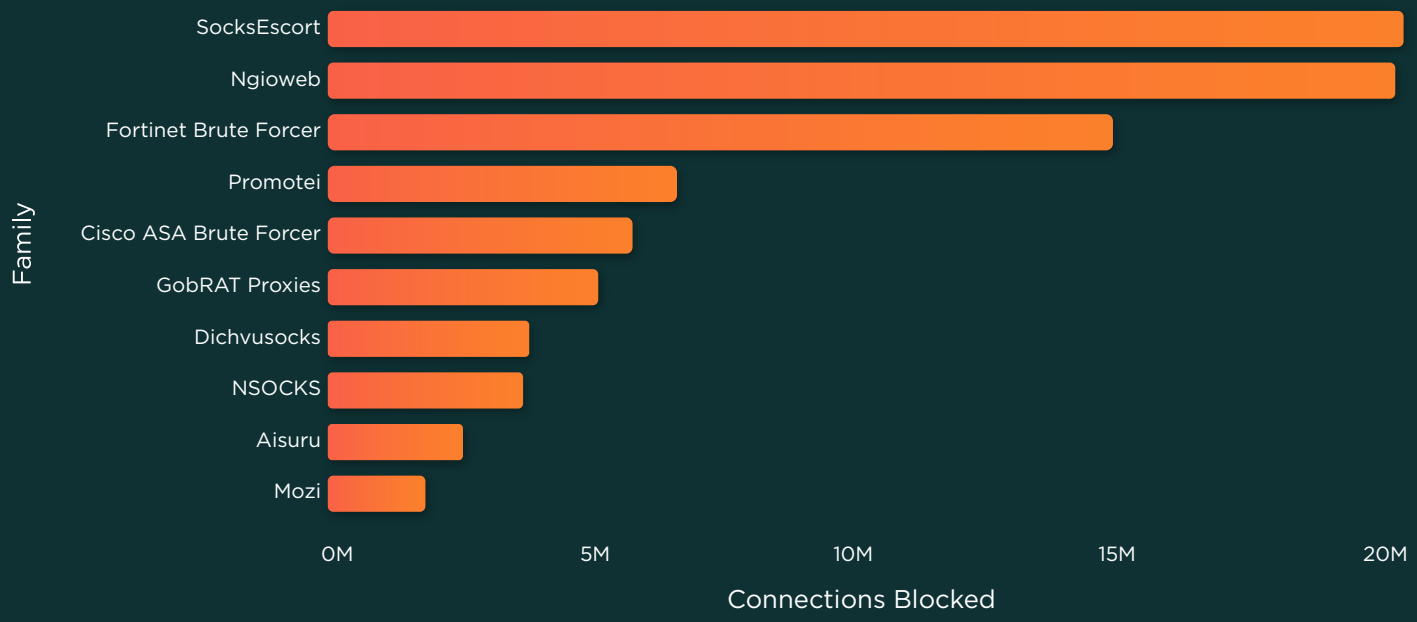

In 2025, Lumen Defender blocked the most threats from malware-backed proxies such as SocksEscort and Ngioweb. IPs employed in these networks provide digital cover for criminal actors to launch attacks while evading detection. Also prominent by volume of blocks in Lumen Defender are from threat actors brute forcing popular edge devices.

Top infostealers based on victim count (2025)

Black Lotus Labs analysis over Lumen's global visibility indicate that Lumma and Rhadamanthys claimed the largest share of victims when compared to other infostealers in 2025.

LUMEN®

Top remote access trojans based on victim count (2025)

Victims of the top RATs are seen through Lumen global visibility. In 2025 there were many more active RATs with more focused targeting, and were not statistically significant overall.

Based on Black Lotus Labs's research, the most frequently blocked threat families in 2025 were botnets and brute forcers. This trend is driven by the widespread presence of malicious proxy services and brute force attack networks, combined with the large number of potential targets like EoL devices and internet-exposed edge devices and services.

While other threat types—such as infostealers, remote access trojans (RATs), and ClickFix pages—were not intercepted as often, the data indicates that these threats are still actively targeting Lumen Defender customers, albeit at a lower volume.

Threats deconstructed

Kimwolf: The distributed-denial-of-service botnet that rose out of Aisuru

Kimwolf emerged in late 2025 as a breakaway operation from Aisuru, which, at the time, was the most powerful distributed-denial-of-service (DDoS) botnet on the internet. Black Lotus Labs tracked Kimwolf as a deliberate infrastructure pivot—executed at speed and fueled by pre-positioned access to residential proxy ecosystems.

Following disruption pressure on Aisuru, Kimwolf's operators quickly rebuilt their control plane. New C2 domains appeared, malware was retooled, and traffic patterns shifted rapidly. Within weeks, the botnet scaled to hundreds of thousands of bots, sustaining massive DDoS capacity while actively evading suppression. As the botnet grew to a critical size, Lumen took action by blocking traffic to and from known C2 nodes across our global network.

The speed and scale of Kimwolf's transition show how future large-scale botnets will evolve under pressure—abandoning, reusing, and regenerating infrastructure faster than defenders can respond.

The crew: Who's behind the operation

Kimwolf is best assessed as a criminally operated DDoS enterprise, rather than a nation-state campaign. While some infrastructure choices and naming conventions drew public attention, the operational behavior points squarely toward profit-driven disruption at internet scale.

Attribution confidence is moderate. Control activity traced through residential Secure Shell (SSH) access, consistent operator interaction patterns, and overlaps with Aisuru infrastructure suggest continuity of personnel rather than a new entrant. Subsequent external reporting aligned with this assessment.

The control room: How the operation was managed

Kimwolf's control architecture reflects a mature operational mindset. Backend C2 nodes act as traffic funnels, aggregating signals from distributed C2s while allowing operators to monitor bot health, redeploy malware, and rapidly shift infrastructure.

Human interaction is evident. SSH sessions originating from residential IP space indicate active management rather than fully autonomous orchestration. When null-routing disrupts a C2, operators typically react within hours—sometimes minutes—standing up replacements and triggering mass malware re-downloads across the botnet. This blend of human-in-the-loop control with automated redeployment allows Kimwolf to scale quickly while remaining tactically flexible.

The infrastructure stack: How the operation was built

Kimwolf relied on a multi-layered architecture:

- **Bots:** Primarily residential and proxy-derived devices, acquired in bulk from compromised proxy services.
- **C2 Layer:** Rapidly rotating domains and IPs, often hosted within small or newly stood-up subnets.
- **Malware Distribution Nodes:** Dedicated servers used to re-seed bots when infrastructure shifted.

Infrastructure lifespans are short by design. C2 nodes are expected to burn quickly, and the operation is built to absorb that loss. When disruption occurs, bots are redirected to new endpoints or instructed to pull updated binaries—demonstrating a logistics-first approach to botnet management.

The routes: How the operation moved across the internet

Kimwolf's defining advantage is its exploitation of residential proxy ecosystems. The operators effectively bulk-purchase bot capacity by exploiting vulnerabilities in proxy services and harvesting devices already positioned behind consumer IP addresses.

Traffic blends seamlessly into residential noise, complicating attribution and making traditional ASN- or geography-based filtering ineffective. The same proxy services that legitimate users rely on for anonymity become Kimwolf's camouflage. This strategy allows Kimwolf to scale faster than botnets reliant on traditional IoT exploitation and regenerate after disruption with alarming speed.

Campaign evolution: How the operation changed over time

- **September 2025:** Aisuru reaches historic DDoS capacity, driven by explosive bot growth.
- **Early October 2025:** Increased disruption pressure reveals architectural seams in Aisuru's infrastructure. New C2 domains and binaries appear.
- **Mid–October 2025:** A distinct botnet identity known as Kimwolf forms, with rapid bot acquisition and infrastructure divergence.
- **October–November 2025:** Repeated null-routing of the Kimwolf botnet triggers faster recovery cycles, tighter infrastructure rotation, and continued scale despite sustained pressure.

Targeting & operational use

Kimwolf is optimized for volumetric DDoS attacks at internet scale. Its targeting is opportunistic rather than sector-specific, consistent with a service-oriented botnet offering disruption capacity rather than bespoke access. It focuses on maximizing concurrent bot availability, sustaining throughput under mitigation pressure, and preserving enough flexibility to redeploy quickly.

Infrastructure positioning suggests that Kimwolf is designed to remain adaptable for future monetization models, including potential resale or integration into broader criminal ecosystems.

Why this operational matters

Kimwolf demonstrates how botnet economics are shifting to prioritize regeneration speed as a primary survival trait, not just stealth. The earliest warning signs to identify Kimwolf were infrastructure choreography like proxy service probing, sudden traffic convergence, and mass malware redeployment rather than attack payloads. These signals are invisible without upstream visibility.

Disruption & defense

Black Lotus Labs's backbone-level visibility allowed defenders to observe Kimwolf's formation in near real time—tracking bot surges, identifying new C2 nodes, and understanding how the operation adapted under pressure.

Through coordinated null-routing, intelligence sharing, and collaboration with industry partners and law enforcement, Lumen disrupted more than 550 Aisuru and Kimwolf C2 nodes in four months. Each disruption forced the operators to expend time, resources, and attention simply to stay operational—underscoring the importance of sustained infrastructure attrition when taking down persistent botnets. Although Kimwolf is no longer active, Black Lotus Labs continues to track the botnet—blocking and updating C2s in Lumen Defender to protect our customers and help keep the internet clean.

Key takeaway

Kimwolf shows that modern DDoS botnets have evolved into fast-moving crews that rebuild as quickly as they are disrupted. The decisive advantage belongs to defenders who can see those crews assembling their infrastructure, exploiting proxy ecosystems, and staging their routes long before the attack begins. In today's threatscape, stopping attacks means watching the neighborhood—not just the vault. Effective defense against vast and flexible DDoS botnets such as Kimwolf requires protecting organization from DDoS traffic and bot activity, as well as disrupting the botnet itself—preventing bots from receiving commands from their C2s through partnered null routes.

Learn more →

LUMEN®

Rapperbot/Aisuru bots detected per day (2025)

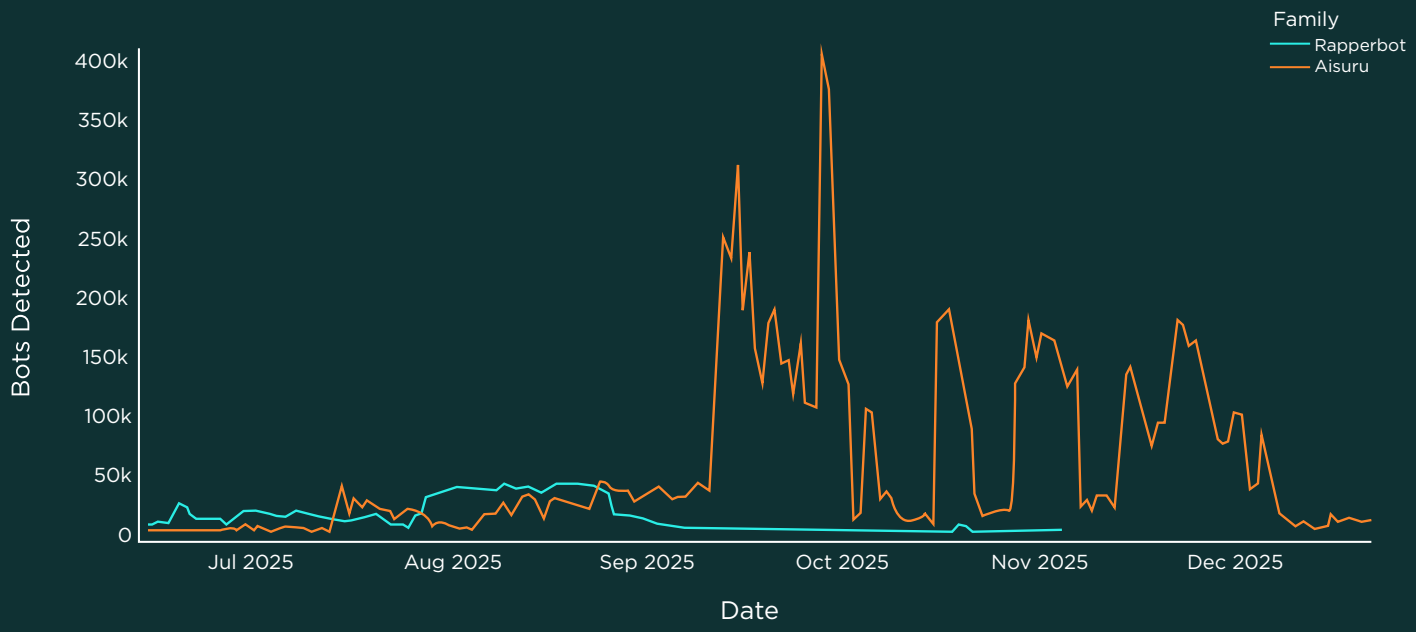

While Rapperbot was the dominant service in the summer of 2025, it competed with Aisuru for control of bots within each other's network—infecting and stealing the other's pools. Following a law enforcement disruption in August, Rapperbot began losing control of its bots—allowing Aisuru to take what remained.

In late September, Aisuru rapidly expanded following the discovery of local-area network exploitation capabilities against large malware-backed proxy networks. This marked the point at which Kimwolf emerged, leveraging the broader reach of devices within the Aisuru botnet.

The peaks and valleys shown in the Aisuru and Kimwolf bot detection image above shows the cycle of continuous discovery, followed by null-routing and regrowth. Black Lotus Labs now sees dozens of 1 Terabits per second (Tbps) an hour.

Kimwolf architecture connecting to proxy service entry gateways (October–November 2025)

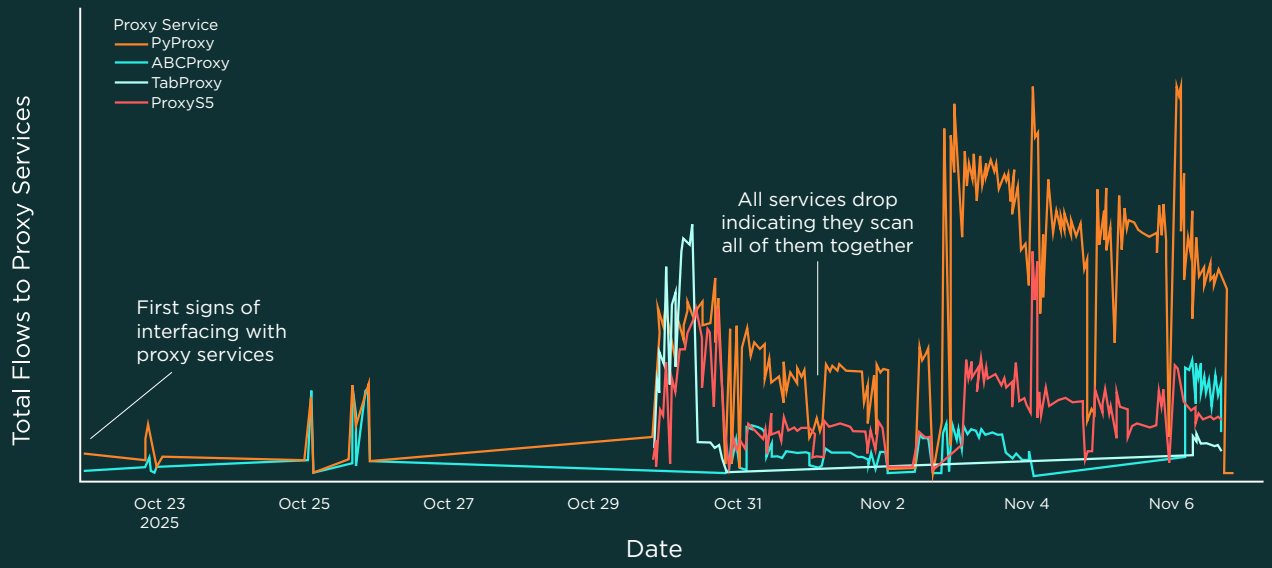

Kimwolf drew its bots from multiple proxies vulnerable to exploits that allowed the threat actors to pivot from proxies into devices on the proxy's LAN to herd them into the bot population, enabling an explosive growth of the DDoS botnet.

The Aisuru DDoS botnet, which emerged in late 2024, caused several major DDoS attacks in 2025. Black Lotus Labs tracked Aisuru until September 2025, when the bot count tripled in just one week.

Soon after, unusual shifts in attack patterns indicated the botnet was directing high volumes of traffic to the service gateways at a handful of dubious proxy services. Black Lotus Labs discovered that Aisuru's 1.8 million bots were generated by exploiting the proxy service and regenerating them there. This evolved botnet, later identified as Kimwolf, launched attacks approaching 30 terabits per second (Tbps).

Having identified the botnet's sources, Black Lotus Labs worked with industry partners and began null-routing the Kimwolf servers as they became active—about 550 in 2025 alone, most of which weren't publicly identified as C2s. Black Lotus Labs has continued this effort into 2026. Throughout our research in tracking this botnet, C2s were continuously updated and added to Lumen Defender's blocking.

The study of this botnet shows that DDoS threats only continue to grow. Not long ago, a 1 Tbps attack was the record. Now, we see attacks up to 30 Tbps. Making matters worse, the malware Kimwolf used to propagate was designed to infect devices within the LAN of its host. Massive DDoS botnets that can multiply quickly will become a feature of the landscape in 2026.

How did Black Lotus Labs discover Kimwolf?

1 Backbone-level traffic observation and anomaly detection

The investigation began with network-wide visibility across Lumen's global backbone. Analysts detected a sharp increase in traffic from bots communicating with known Aisuru C2 nodes in September 2025 via large-scale traffic analysis, allowing them to quantify rapid bot growth and identify backend infrastructure patterns.

2 Identification of backend C2 aggregation points

Black Lotus Labs identified a backend C2 server that acted as a traffic funnel for multiple known Aisuru C2 nodes. By observing how traffic converged on this backend node, we inferred the centralized control infrastructure behind the botnet.

3 Traffic-driven infrastructure pivoting

From the backend C2, our analysts pivoted based on observed traffic behavior, identifying:

- Residential SSH connections
- Repeated access patterns
- Domain and IP reuse across infrastructure

This traffic analysis led us to infrastructure tied to residential proxy services, rather than traditional VPS-only botnet hosting.

4 Correlation of IPs, domains, and hosting providers

Our methodology relied heavily on correlating IP addresses, domain names, and hosting providers over time. Analysts tracked:

- Domain name changes
- IP reassignments
- Hosting provider concentration (e.g., Resi Rack LLC)

This correlation showed that newly observed domains and IPs were operationally linked to earlier Aisuru infrastructure, even when names and binaries changed.

5 Binary retrieval and comparative analysis

When infrastructure exposed downloadable binaries, analysts:

- Retrieved malware binaries from open ports
- Re-queried the same infrastructure hours later
- Compared changes in binaries and callback behavior

This revealed rapid infrastructure evolution, with binaries shifting from pointing to Aisuru C2s to newly observed Kimwolf C2 domains, indicating the emergence of a distinct but related botnet.

Rhadamanthys: A malware-as-a-service platform built like a startup and run like a crime syndicate

Rhadamanthys is a commercialized operation designed to scale access, automate abuse, and lower the barrier to entry for cybercrime. Since its emergence in late 2022, Rhadamanthys has grown into the largest information-stealer platform by volume, evolving rapidly from a niche infostealer into a full-fledged malware-as-a-service (MaaS) ecosystem.

Black Lotus Labs tracked Rhadamanthys's expansion through Lumen global NetFlow telemetry and observed a steady rise in victims throughout 2024 and 2025—including a sharp acceleration following the disruption of competing infostealers. By October 2025, Rhadamanthys had crossed a critical threshold, impacting over 12,000 victims globally and confirming its position as a dominant force in the malware economy.

The crew: Who's behind the operation

Rhadamanthys is best understood as a criminal service provider. Its operators handle platform management, distribution enablement, and customer acquisition, while affiliates exploit victims downstream.

The operation is economically driven—focused on adoption, retention, and continuous product refinement. This reflects the broader specialization of cybercrime, where access, infrastructure, and exploitation are separated across actors. Attribution remains intentionally diffuse, helping Rhadamanthys absorb law enforcement pressure while sustaining operations through affiliates.

The control room: How the operation was managed

Rhadamanthys operates with a level of professionalization uncommon even among mature criminal platforms. The operators maintain a curated sales portal branded as "RHAD Security," complete with customer support, subscription tiers, and ongoing feature updates. This approach attracts a broad range of users by projecting an image of credibility and reliability within the cybercriminal ecosystem.

From an infrastructure standpoint, customers can either communicate directly with victims via the shared C2 framework or route traffic through an integrated reverse-proxy layer (Elysium) to further obscure attribution. Including proxy infrastructure, Black Lotus Labs tracked a daily average of 300 active servers—with a peak of 535 servers in October 2025.

Infostealers detected per day (2025)

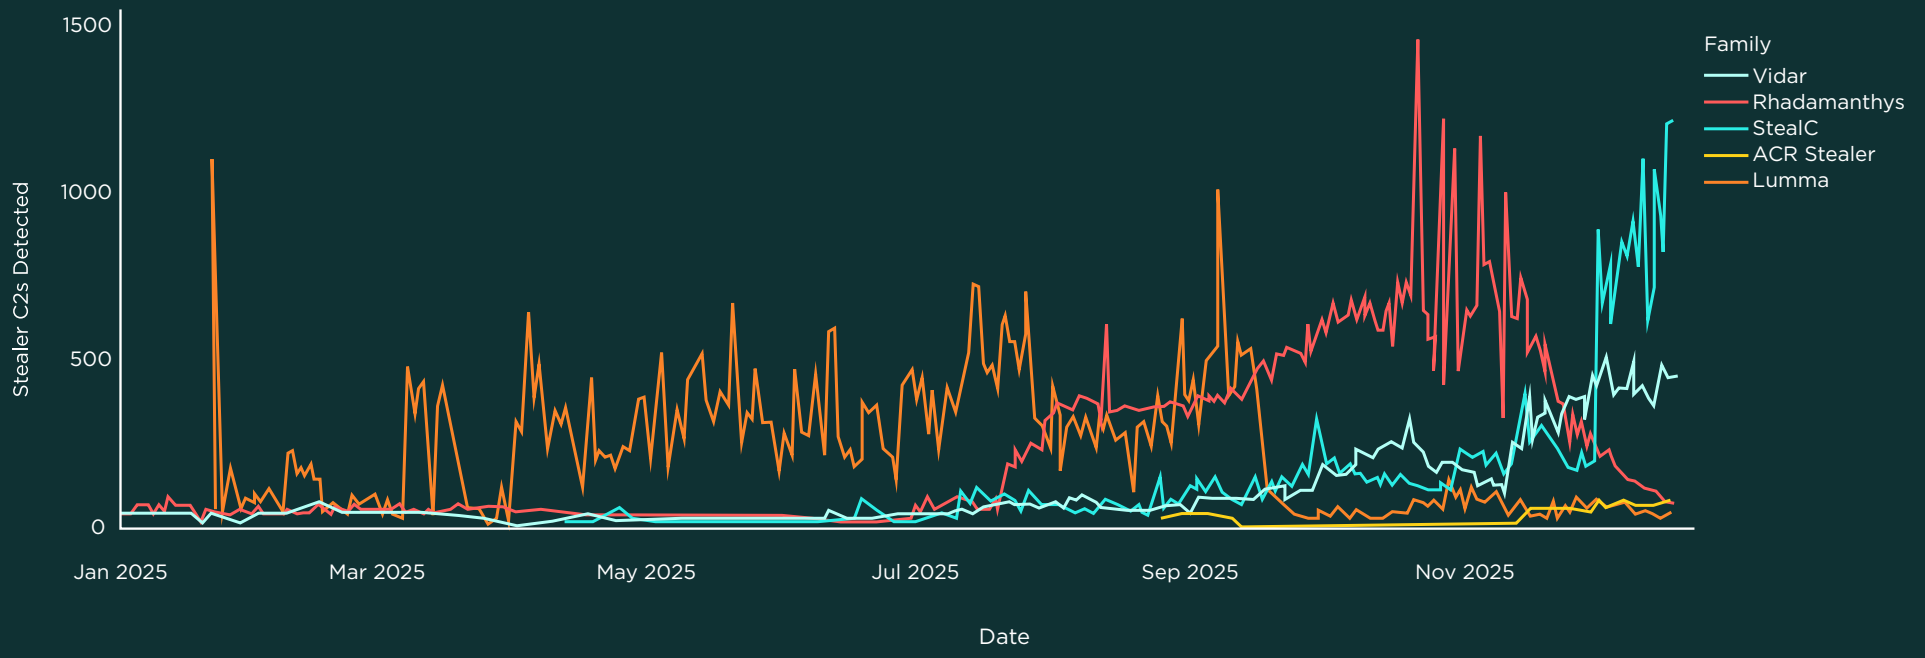

This graph demonstrates the impact of Black Lotus Labs collaborative efforts with law enforcement, as well as the dynamic competition among infostealer malware.

The Lumma Infostealer was disrupted in November 2025, with Black Lotus Labs assisting Microsoft's efforts. As its C2 activity sharply declined, Black Lotus Labs detected increased activity from the remaining stealers in the marketplace. Rhadamanthys quickly rose to prominence. Black Lotus Labs focused on the new threat and shared its telemetry with international law enforcement.

When Rhadamanthys was disrupted during Operation Endgame, other stealers such as Vidar and StealC surged—with StealC now leading in C2 detections. The evolving market share among these threat actors signals ongoing shifts to monitor in 2026.

Throughout our research in tracking these infostealers, indicators of compromise (IOCs) were continuously updated and added to Lumen Defender's blocking to protect our customers in near-real time.

LUMEN®

The infrastructure stack: How the operation was built

Rhadamanthys relied on a modular, layered architecture optimized for stealth and resilience. Its infrastructure was globally distributed, with over 60% of C2 servers hosted in the United States, Germany, the United Kingdom, and the Netherlands.

Rhadamanthys's malware consists of two main components: a loader and an exfiltration module for extracting collected credentials. The malware is continuously refined for stealth using complex anti-analysis techniques, tricky installation methods, and AI to extract target data.

Crucially, over 60% of Rhadamanthys C2s showed zero detection on VirusTotal at the time of Black Lotus Labs's original reporting—highlighting how effectively the platform evaded traditional reputation-based defenses.

The routes: How the operation moved across the internet

Rhadamanthys's delivery strategy evolved rapidly over the course of the operation. Early campaigns relied heavily on malicious Google advertisements. However, as adoption grew, distribution shifted toward forum-driven referrals, affiliate sharing, and curated access.

The platform's optional proxy layer allows operators and customers to decouple delivery infrastructure from command infrastructure, making network defense more difficult and attribution slower.

By routing victim communications through intermediary nodes, Rhadamanthys blends malicious traffic into ordinary hosting and cloud patterns—a recurring theme across modern MaaS ecosystems.

Campaign evolution: How the operation changed over time

Rhadamanthys's growth has been opportunistic and adaptive:

- **Late 2022–2023:** Initial emergence and validation of the infostealer.
- **2024:** Expansion of features, stealth improvements, and early MaaS branding.
- **2025:** Rapid scale-up following the disruption of competitor platforms, culminating in record victim counts.

Short-term slowdowns, including a brief dip in September 2025, were quickly followed by infrastructure expansion and renewed propagation—indicating deliberate operational pacing and persistence rather than decline.

Targeting & operational use

Rhadamanthys casts a wide net, with infections observed across multiple regions and industries. The highest victim concentrations were observed in Brazil, Argentina, Turkey, India, and the United States. The platform is primarily used for large-scale credential theft, session and identity harvesting, and establishing initial access for downstream activity, including ransomware and data theft operations.

Why This operation matters

Rhadamanthys represents the continuation of a defining threat trend: cybercrime platforms now scale like SaaS companies. It shows us how cybercriminals are professionalizing operations to new heights, creating enterprise-grade platforms complete with subscription access, feature roadmaps, built-in obfuscation, and rapid adaptation to market disruptions.

Disruption & defense

Black Lotus Labs supported Operation Endgame by tracking Rhadamanthys's infrastructure at scale, sharing intelligence with international law enforcement, and proactively blocking newly observed C2 servers across Lumen global backbone. IoCs were continuously fed into Lumen Defender to protect customers from emerging infrastructure, demonstrating how backbone-level visibility can compress the time between observation and disruption.

Of the roughly 200 daily Rhadamanthys C2 servers that Black Lotus Labs tracked, industry tracked approximately 20%. This is largely due to the difficulty in tracking this MaaS platform. Black Lotus Labs tracks Rhadamanthys through a combination of network attributes in Transport Layer Security (TLS) certificates and netflow. Black Lotus Labs netflow holdings highlight additional certificate patterns discovered on C2s, which call out to a backend Rhadamanthys server. These netflow connections highlight additional certificates of interest. With the additional 80% novel C2 and associated netflow coverage, Black Lotus Labs can illuminate as many as 1000 to 3000 unique victims per day.

Key takeaway

Rhadamanthys proves that modern malware can be marketed, sold, and supported like any other SaaS service. To block it, defenders need to do more than just react quickly to new playloads. They must leverage network-level intelligence to identify when a criminal platform is scaling its infrastructure and intervene before the next wave of customers can ever press "deploy." Enabled by Black Lotus Labs' threat intelligence, Lumen security products automatically block updated Rhadamanthys C2s as Black Lotus Labs sees them.

Learn more →

LUMEN®

Brute force attacks: A classic heist trick, rebuilt for scale

Brute force attacks are often dismissed as unsophisticated—a blunt instrument in a world of advanced malware and 0-days. But at internet scale, they are anything but crude.

Threat actors have access to billions of stolen credentials on the dark web. Brute force attacks make use of these identity factors, PoC exploits, and AI-driven tools to analyze leaked datasets and target internet-exposed devices and services. Modern brute force campaigns are carefully staged operations, built on automation, distributed infrastructure, and precise targeting of exposed services at the network edge. Long before an account is compromised or a session is hijacked, attackers assemble botnets, credential lists, and attack infrastructure designed to sustain pressure over time.

Through analytics applied across Lumen global NetFlow, Black Lotus Labs has observed a steady rise in both the frequency and intensity of brute force activity over the past several years.

The crew: Who's behind the operation

Brute force attacks are not owned by a single threat actor. They are a shared tactic used across criminal, hybrid, and nation-state operations.

Criminal groups rely on brute force for credential harvesting and initial access. Ransomware operators use it to validate leaked credentials at scale. Nation-state actors, including those from Iran and Russia, have repeatedly leveraged brute force techniques to gain footholds in strategically valuable networks.

While any individual brute force attempt normally has a low success rate, a bot-distributed attack increases its strength through volume against thousands of targets simultaneously.

The control room: How the operation was managed

Modern Brute Force campaigns are machine-driven. Attackers automate credential rotation, target enumeration, retry logic, and infrastructure regeneration. Rather than randomly spraying the internet, many campaigns focus on specific services or device classes—adapting in real time as defenders respond. This tempo enables threat actors to apply high volume, repetitive pressure across thousands of targets simultaneously.

The infrastructure stack: How the operation was built

Brute force attacks rely on scale, not stealth. Botnets (often the same ones powering proxy networks or malware delivery) provide the computational leverage needed to sustain high-volume attempts.

Over the past five years, the growth of large botnets has dramatically increased the effectiveness of brute force campaigns. Microsoft reported a global peak of 11,000 password-based attacks per second in April 2023—a number that reflects widespread infrastructure availability.

The routes: How the operation moved across the internet

Brute force campaigns overwhelmingly target the network edge. These services sit at the intersection of convenience and exposure. Attackers don't need to breach the network—they simply need to find a service that accepts authentication attempts.

Campaign evolution: How the operation changed over time

The widespread availability of leaked credentials has transformed brute force from blind guessing into credential validation at scale. Increasingly, these attacks also precede the disclosure of new vulnerabilities, suggesting attackers probe authentication paths while waiting for exploit details to surface.

Targeting & operational use

In 2022, over 80% of breaches against web applications and internet-exposed services involved brute force or stolen credentials. These attacks are foundational. They enable ransomware deployment, establish long-term access, and support espionage and reconnaissance. Even smaller scale threats can have outsized impact when directed at the right service at the right moment.

Why this operation matters

Brute force attacks are a reminder that modern cyber operations don't always start with malware. The earliest signals often appear as repetitive, high-volume authentication attempts—visible only when defenders look beyond individual logs and examine network-wide behavior. Endpoint alerts catch the failure. Backbone telemetry reveals the preparation.

Disruption & defense

To counter this threat, Black Lotus Labs Brute force detections based on backbone visibility and deployed in Lumen Defender continuously analyze Lumen global NetFlow to detect activity in near-real time. By monitoring edge services globally, evaluating every IP communicating with high-risk services under a Zero Trust model, and tracking how attacker infrastructure shifts in response to blocking, Black Lotus Labs can identify emerging campaigns and infer trends across similar enterprise targets. We also monitor sustained campaigns and block malicious traffic for our customers.

This intelligence feeds directly into our Lumen Defender security solutions, enabling automated mitigation and decreasing customers' workload and alerts. Most brute force campaigns are sustained against specific devices and services over the course of several days. Lumen Defender security solutions leverage Black Lotus Labs' intelligence to protect customers using targeted devices or services from the first day that a new wave of brute force attacks is detected until the campaign concludes.

To date, Black Lotus Labs has null-routed traffic from hundreds of consistently malicious IPs across the Lumen global backbone, reducing attack impact for customers and the internet at large.

Key takeaway

Brute force attacks have evolved. They are no longer random noise, but a coordinated rehearsal for larger operations at scale. To stop them, defenders must look beyond failed login attempts into the internet-exposed services and EoL devices that are vulnerable to opportunistic attackers.

Learn more →

LUMEN®

SystemBC: The high-bandwidth proxy crew built for volume

Black Lotus Labs uncovered new infrastructure behind the SystemBC botnet in September 2025. This proxy-producing network is composed of 80+ C2 servers and a daily average of 1,500 victims. Unlike the typical proxy botnet that hides inside residential routers and IoT devices, SystemBC is built primarily of compromised VPS systems sourced from several large commercial providers.

These VPS-based proxies can push massive throughput for long periods without the same fragility and user disruption that would expose a residential proxy pool. SystemBC's model focuses on utility rather than discretion—delivering high-volume proxy capacity that criminal ecosystems can plug into on demand.

The crew: Who's behind the operation

SystemBC is a commoditized criminal capability. It was originally documented in 2019, sold and discussed in underground forums, and repeatedly adopted by multiple threat groups over time. Black Lotus Labs has observed its ecosystem ties to major criminal tooling and activity patterns, including historical use alongside entities like IcedID and Trickbot. SystemBC also caught the attention of law enforcement during Operation Endgame due to its role in enabling ransomware delivery pipelines.

Intent is consistently criminal and infrastructural. SystemBC seeks to build and maintain a proxy layer that other groups can use to move malicious traffic at scale, conduct scanning and exploitation, brute force credentials, and support downstream stages of ransomware-adjacent operations.

The control room: How the operation was managed

SystemBS's "control room" focuses on distribution and resale. Operators run more than 80 C2s and appear to concentrate them under a single Autonomous System (AS)—a footprint small enough that those C2s represent roughly 10% of the AS's network. Users of the proxy network connect to SystemBC's C2s over high-numbered ports, then get routed onward through to victim proxies.

Crucially, SystemBC's infrastructure is not exclusive. Black Lotus Labs observed these same bots sold into multiple downstream services, including at least two Russia-based proxy services, one Vietnamese proxy service, and a Russian parsing/web-scraping service.

The infrastructure stack: How the operation was built

SystemBC's victims are the operation's core asset. While Black Lotus Labs did not determine SystemBC's initial access vector, the compromised VPS servers tended to be highly vulnerable. On average, each victim showed 20 unpatched common vulnerabilities and exposures (CVEs) and at least one critical CVE, with one observed server showing 160+ unpatched vulnerabilities.

Telemetry also pointed to a key recruiting and hosting node. 104.250.164[.]214 appeared to be both the source of recruitment attacks and the sole host of SystemBC malware samples. Initial access attempts are sent over port 443. If successful, the victim calls back on port 80 to download a shell script (with Russian comments) that forces the download of 180+ samples of SystemBC malware and runs them simultaneously under different filenames.

The routes: How the operation moved across the internet

SystemBC's getaway routes are built for throughput rather than disguise. Where most malware-backed proxy botnets throttle traffic to avoid burning residential IPs or getting listed, SystemBC takes a different approach. Black Lotus Labs observed a single proxy IP generating 16+ GB of proxy traffic in a 24-hour window—significantly higher than typical malware-based proxy networks.

As a result, evasion is not the selling point. Nearly 100% of bots eventually show up on block list sites due to mass scanning, exploitation, and brute force activity. SystemBC's ecosystem treats this as the cost of doing business rather than failure. When an IP burns out, there are more vulnerable VPS hosts to replace it.

Campaign evolution: How the operation changed over time

SystemBC had staying power across multiple years because it evolved into something more useful than just a botnet. It functions as a proxy factory that can be incorporated into multiple criminal services at once and supports an economy where operators and resellers can segment proxies by "quality" and detection risk.

It's also become a building block for tiered offerings that span the full pipeline, from brute force and reconnaissance to targeted operations. That evolution becomes clearest when examining the downstream services that rely on it.

Targeting & operational use

SystemBC primarily targeted VPS hosts due to the large number of easy-to-exploit, poorly maintained servers. Once compromised, those servers become high-bandwidth proxy endpoints for various criminal workflows. Black Lotus Labs observed several major usage patterns that show how SystemBC infrastructure is repurposed downstream by multiple actors, each optimizing for different stages of the attack lifecycle:

- The operators themselves used the proxy network for WordPress credential brute forcing, likely feeding harvested credentials into broker ecosystems that monetize access (including injection of malicious code into sites).
- A large Russian web-scraping/parsing service integrated SystemBC proxies into its infrastructure, likely selling them as "regular" proxies to customers who don't care about blacklist status.
- VN5Socks (also known as Shopsocks5) offered these IPs, apparently diversifying supply after disruption in other proxy ecosystems.
- REM Proxy functioned as the primary force multiplier for SystemBC, transforming its high-bandwidth VPS botnet into a tiered, ransomware-ready proxy service across the full attack lifecycle—from phishing and credential harvesting to data exfiltration and ransomware operations by groups such as AvosLocker and Morpheus.

Why this operation matters

Instead of chasing stealth through residential disguise, SystemBC chased capacity—turning vulnerable VPS fleets into a renewable pool of bandwidth. That bandwidth then powers scanning waves, brute force at scale, credential harvesting, spam distribution, and the scaffolding that ransomware ecosystems and access brokers depend on. The emergence of high-volume proxy infrastructure that converts general internet exposure (unpatched servers) into criminal advantage was key for identifying, tracking, and eventually disrupting SystemBC's operation.

Disruption & defense

Throughout our investigation of SystemBC, Black Lotus Labs added IoCs into Lumen Defender to deny traffic to or from our customers to SystemBC's C2 infrastructure as it evolved. At the time of publication, Lumen blocked all traffic to or from SystemBC and REM Proxy and released IoCs to help the community detect and disrupt the operation. Our research effort also included collaboration with industry partners, including Spur and Infoblox.

Key takeaway

SystemBC shows what happens when cybercrime stops hiding behind home routers and starts renting bandwidth power instead. By turning vulnerable VPS infrastructure into disposable, high-bandwidth proxy capacity, criminals can flood the internet with scanning, brute force, and delivery traffic before replacing what got burned.

Learn more →

LUMEN®

DanaBot: The malware delivery crew that ran like a franchise

First seen in 2018 as a banking trojan, DanaBot evolved into a flexible toolkit for information theft and establishing initial access that can enable downstream activity, including follow-on malware delivery (such as Latrodectus) and ransomware-adjacent access workflows. It remained highly operational until May 2025, when it took a major hit during Operation Endgame II. This coordinated disruption effort was supported by Black Lotus Labs and Team Cymru alongside industry peers and law enforcement.

Following Operation Endgame II, DanaBot resurfaced in November 2025 with "Version 669"—leveraging complex multi-stage attacks to target financial institutions, cryptocurrency wallets, and individual victims.[5] This resurgence was first reported by Zscaler ThreatLabz, underscoring DanaBot's ability to reconstitute after disruption. While this re-emergence is notable, the analysis in this report focuses primarily on DanaBot's infrastructure, operational model, and lifecycle as observed by Black Lotus Labs and Team Cymru during the period leading up to the May 2025 takedown—when its architecture, scale, and affiliate ecosystem were most fully exposed.

Ultimately, the professionalization of its infrastructure is what sets DanaBot apart. Its large, persistent, quietly operating C2 footprint was structured like a service and designed to keep the true operators insulated while affiliates worked the street-level infections.

The crew: Who's behind the operation

DanaBot is best understood as a criminal MaaS ecosystem rather than a single cohesive actor. Black Lotus Labs and Team Cymru assess the platform is segmented among multiple users (affiliates) who purchase access and operate in partially siloed infrastructure "lanes" depending on their subscription level and needs.

Intent is therefore mixed but consistently criminal—credential theft and data collection on one end, and access enablement on the other. Infections can quickly pivot into additional malware delivery or handoff to other actors for follow-on monetization.

The control room: How the operation was managed

DanaBot's operational advantage came from how it managed control at scale without exposing the true management layer.

At the front line, DanaBot maintained an average of nearly 150 active C2 servers per day, with roughly 1,000 daily victims across 40+ countries—one of the largest platforms by C2 count active in 2025. Yet the operation stayed unusually stealthy. Only 25% of its C2 infrastructure had a VirusTotal detection score greater than zero, suggesting that a significant portion of its infrastructure remained undetected. This is likely due to selecting fewer targets than other loaders of its kind, as well as cycling operations around high-profile events.

Behind that scale was a multi-tier proxying model. Victims talk to Tier 1 C2s, which relay to Tier 2 servers, which then relay to Tier 3 infrastructure, where panels are believed to be hosted. This layered relay design (similar in concept to obfuscation patterns seen in other botnet ecosystems) kept the real operators buffered behind multiple curtains.

The infrastructure stack: How the operation was built

DanaBot's architecture is a logistics chain:

- **Victim → Tier 1 (T1) C2s:** Infected systems beacon to one or more T1 nodes over TCP/443.
- **T1 → Tier 2 (T2) C2s:** T1 nodes are controlled by one of several T2 C2s.
- **T2 → Tier 3 (T3) infrastructure:** T2s relay upstream to T3 nodes (observed as Russian IP space), which appear to represent the final tier where panels likely live.
- **T3 → backup behavior:** At least two identified T3s were observed transferring large volumes of data monthly to the same server over TCP/2048—behavior consistent with backup server activity.

The platform's "service model" is visible in its segmentation. At any given time, at least five to six T2 servers were active, with evidence that some affiliates had dedicated T2 capacity while smaller affiliates likely shared.

Infrastructure longevity was also notable. Average C2 lifetimes extended beyond a month, with 25% lasting more than two months. However, these C2s still avoided widespread reputation scoring—reinforcing DanaBot's preference for quieter, more targeted workflows over noisy mass campaigns.

The routes: How the operation moved across the internet

DanaBot's getaway routes relied on layered relays and selective concealment. A portion of victim data could be routed through Tor, meaning the visible bot population is likely an undercount. This also made victim enumeration and impact assessment more complex for defenders observing only direct C2 relationships.

The platform also appeared to "play the calendar." Black Lotus Labs and Team Cymru observed infrastructure surges and pauses that suggest deliberate timing around high-attention events—such as a significant increase in C2 volume leading up to the November 2024 U.S. election, followed by a lull, then ramping again through December 2024 holidays. That pattern is consistent with actors exploiting predictable peaks in user activity and lures, achieving maximum yield with minimal noise.

Campaign evolution: How the operation changed over time

DanaBot's evolution demonstrates continuous refinement and persistence:

- **2018:** Emerges as a banking trojan.
- **2024–2025:** Operates as a mature infostealer and access enablement platform with stable upstream/back-end patterns since June 2024, while varying T1 footprints and activity rhythms.
- **April 2025:** A notable period where much of one "cloud" cluster went dark for weeks before reappearing—possibly operational pause or infrastructure updates.
- **May 2025:** Still highly active with hundreds of distinct C2 IPs observed, until disruption under Operation Endgame II.
- **2026:** In November 2025, Cyber Security News, informed by research from Zscaler ThreatLabz, reported that DanaBot made a significant comeback with its new version 669 after a period of inactivity following Operation Endgame II. As of the date of this reporting, the MaaS ecosystem is still active.

Rather than constant reinvention, DanaBot's evolution shows operational maturity: insulation, modularity, affiliate scaling, and the ability to surge opportunistically while maintaining a low reputation footprint.

Targeting & operational use

DanaBot victims appeared across 40+ countries, with Brazil, Mexico, and the United States consistently among the most impacted. The victim pool was primarily the residential IP space, but higher-value infections were also observed, including law firms and universities.

Infection duration was often brief. Half of victims only communicated with the DanaBot C2 for a single day, and 75% of infections lasted less than three days. This suggests many affiliates used DanaBot as a fast-hit collection and staging tool, quickly transitioning to downstream actions once value was extracted.

Why this operation matters

DanaBot shows how "industrialized" cybercrime is evolving in 2026. Rather than executing a single malware campaign, its platform approach leverages a large infrastructure footprint, quiet operations, and a highly disciplined business model to distribute risk across affiliates.

It also reinforces the idea that a platform can appear modest in terms of victim counts while still maintaining a massive, stealthy infrastructure layer. This infrastructure layer is critical for understanding how operations are staged, siloed, timed, and scaled.

Disruption & defense

DanaBot's disruption underscores the value of ecosystem-level collaboration. Black Lotus Labs and Team Cymru supported Operation Endgame II by mapping DanaBot's infrastructure and contributing insight to the broader community effort alongside law enforcement and industry partners.

Following the international law enforcement disruption of DanaBot, its operators began to rebuild with "Version 669." As Black Lotus Labs tracks this resurgence, we have added IoCs from DanaBot's new network into Lumen Defender to protect our customers from malicious traffic from new and existing C2s. As security researchers continue monitoring DanaBot's new network with "Version 669," a similar level of cooperation and information-sharing will be necessary to ensure that companies and individual users alike are protected from this persistent MaaS platform.

Key takeaway

DanaBot proves that modern cybercrime is becoming a service business. Affiliates, tiers, and insulation layers keep the real operators offstage while the platform quietly moves access and data at scale. Even when victim counts look manageable, a large, low-detection infrastructure footprint can signal a far more serious reality.

Learn more →

LUMEN®

5socks botnet: The long-running proxy heist hiding in plain sight

The 5socks botnet (originally reported by Black Lotus Labs as Classic Rock) doesn't rely on 0-days, custom malware frameworks, or rapid-fire campaigns. Instead, it focuses on stealing value through patience rather than speed.

For over a year, Black Lotus Labs tracked this criminal proxy network as it quietly infected thousands of IoT and EoL devices, converting them into residential-grade proxies for hire. Through backbone-level visibility, researchers observed a weekly average of 1,000 unique active bots communicating with C2 infrastructure hosted in Turkey, with victims spread across more than 80 countries.

5socks's longevity and discipline set this botnet apart. According to its own website, the service claimed to have operated since 2004, suggesting an operation refined over decades to avoid attention, abuse outdated technology, and blend seamlessly into everyday internet traffic.

The crew: Who's behind the operation

5socks is assessed as a criminal proxy operation, not tied to a nation-state or known APT cluster. There is no evidence of ideological motivation or targeted espionage. Instead, the botnet serves as foundational anonymity infrastructure. Think of 5socks as a utility layer that other criminals rely on to conduct fraud, brute force attacks, DDoS activity, and data exploitation while hiding behind legitimate-looking residential IP addresses.

The control room: How the operation was managed

5socks's control model is intentionally minimal. The botnet's operators manage a small cluster of five C2 servers, four of which communicate with infected devices over port 80, while a fifth receives data via UDP on port 1443—likely used for data storage. There's no evidence of complex dashboards or orchestration layers. Instead, the system favors simplicity and durability.

Critically, the proxy service operates on an open-access model. Proxies require no authentication, meaning anyone who discovers an IP and port (whether a paying customer or an opportunistic threat actor) can use it. This design choice dramatically expands the blast radius of the operation and allows activity far beyond what the operators may directly monetize.

The infrastructure stack: How the operation was built

5socks is powered by compromised IoT and SOHO devices, almost exclusively in the residential IP space. These devices are typically EoL, unpatched, sit outside active monitoring, and are cheap and easy to re-exploit.

Black Lotus Labs does not assess the use of 0-day or 1-day vulnerabilities. Instead, the operators rely on years-old exploits that remain effective because the target population cannot be updated or secured. This strategy allows 5socks to maintain an average bot lifecycle of over a week while constantly replenishing new devices. 5socks is also effective at evading detection. Only 10% of proxies appeared malicious in VirusTotal at the time of our original reporting.

The routes: How the operation moved across the internet

By operating entirely within residential IP ranges, 5socks allows downstream attackers to blend into legitimate user traffic, bypassing geofencing, ASN-based blocking, and IP reputation controls. From a defender's perspective, traffic appears to come from ordinary homes rather than an obviously hostile infrastructure. This allows criminals to disguise their malicious activity behind seemingly innocuous IoT and SOHO devices.

Campaign evolution: How the operation changed over time

Unlike newer proxy services that chase scale aggressively, 5socks favors stability and survival. Lumen telemetry shows a weekly average of 1,000 active bots, victims across 80+ countries, and heavy concentration in the United States, followed by Ecuador and Canada. The botnet appears deliberately under-advertised relative to competitors like CloudRouter or Proxy.AM, suggesting the operators prioritize not drawing attention over rapid growth. This restraint has likely contributed to its long operational lifespan.

Targeting & operational use

5socks's customers are the ones choosing its targets. Observed use cases include ad fraud, brute force and credential stuffing, DDoS attacks, and exploitation of victim data. Because proxies are open-access and residential, they are routinely abused by multiple actors simultaneously. The same infrastructure can support competing criminal campaigns at the same time, often without the operators' direct involvement. This makes attribution downstream nearly impossible.

Why this operation matters

5socks reinforces a central truth of the 2026 threatscape: anonymity infrastructure is now the backbone of cybercrime. As long as millions of EoL devices remain online, proxy botnets like this will continue to thrive. Endpoint security, IP reputation, and geolocation controls struggle against threats that look indistinguishable from legitimate users. For CISOs and other security leaders, this underscores the importance of timely patching and rigorous device management.

Disruption & defense

Through Lumen global backbone visibility, Black Lotus Labs mapped 5socks's architecture over time and collaborated with the U.S. Department of Justice, the FBI, and the Dutch National Police. We disrupted the operation by null-routing traffic to and from known control points across its global network and shared IoCs with the broader security community. Spur also contributed to the investigation. This effort highlights how infrastructure-level insight enables defenders to cut off criminal supply chains, not just individual attacks.

Key takeaway

5socks proves that cybercrime doesn't always rely on cutting-edge exploits. Sometimes, it just needs patience and a steady supply of vulnerable, outdated hardware. By turning forgotten devices into anonymous proxy routes, criminals can operate for years in plain sight. Defenders who want to stop the threat must look further into the infrastructure criminals use to blend in before the actual attack ever begins.

Learn more →

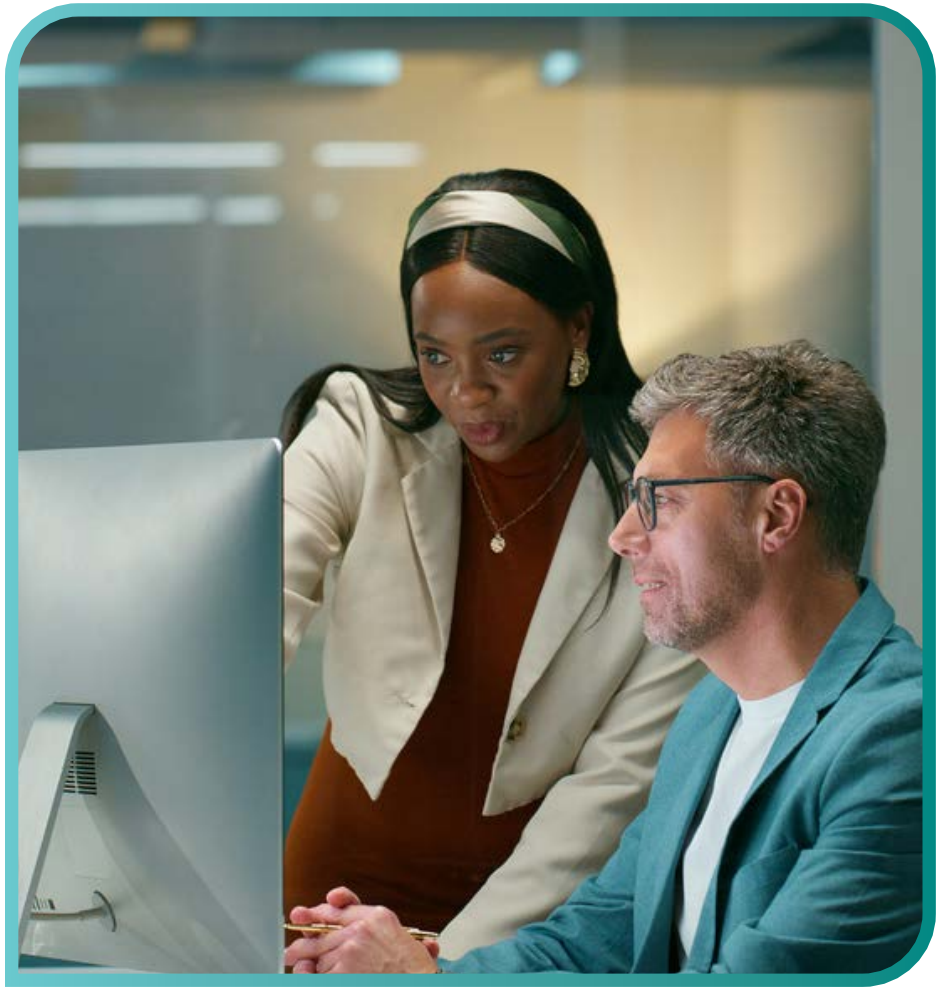

LUMEN®

J-magic campaign: The invisible backdoor living on routers

J-magic backdoor attacks are the silent stakeout of the cybercrime world. Rather than scanning aggressively or beaconing outward, the operators behind J-magic planted a passive listener directly onto enterprise-grade Juniper routers.

The operation appears to have begun in mid-2023, with the earliest known sample uploaded in September 2023, and remained active into at least mid-2024. During that time, the attackers used carefully crafted "magic packets" embedded in otherwise normal TCP traffic to selectively wake infected routers and establish control. Until the right signal arrived, the infrastructure stayed dark—giving defenders nothing to alert on and no obvious indication that an attack was even being planned.

The crew: Who's behind the operation

Black Lotus Labs tracks J-magic as unaffiliated with known public threat clusters, despite some low-confidence technical overlap with prior cd00r-based campaigns such as SeaSpy. The intent appears strategic rather than opportunistic. Targeting enterprise Juniper routers—many of which were configured as VPN gateways—suggests an actor focused on quiet access, credential interception, and network positioning rather than immediate disruption or monetization.

Notably, the inclusion of an RSA-based challenge mechanism strongly suggests the operators were actively defending their own infrastructure—preventing other threat actors from discovering and repurposing the backdoor as seen with other nation-state actors like Secret Blizzard.

The control room: How the operation was managed

J-magic has no traditional control plane. Instead, the "control room" exists only when the operator chooses to reveal it by sending a magic packet that satisfies one of five precise conditions embedded in TCP headers or payloads. Once triggered, the router initiates a reverse shell to a callback IP and port specified inside the packet itself. The operator must then pass a cryptographic challenge using a hard-coded RSA public key before being granted shell access.

This design minimizes coordination overhead while maintaining absolute control. Only someone with the correct packet format and cryptographic response can interact with the implant.

The infrastructure stack: How the operation was built

J-magic is built to live where defenders least expect malware to reside. The implant is a custom variant of cd00r, executed entirely in memory, with no firmware modification and no persistent disk artifacts. Upon execution, it renames itself to mimic a legitimate Junos OS process and overwrites its command-line arguments to erase forensic clues.

Using eBPF, the malware passively inspects all inbound TCP traffic on a specified interface and port. If a packet matches one of the five predefined J-magic conditions, the malware forks a child process and establishes an encrypted reverse shell.

This architecture is designed to exploit the lack of host-based monitoring on enterprise routers. These devices are rarely power-cycled. Malware tailored for routers is designed to take advantage of long uptime and live exclusively in-memory, allowing for low-detection and long-term access compared to malware that burrows into the firmware. Routers on the edge of the corporate network or serving as the VPN gateway are the richest targets, opening avenues to the rest of a corporate network.

The routes: How the operation moved across the internet

J-magic packets are embedded inside legitimate TCP traffic, often originating from public VPN or proxy services, allowing the trigger signal to blend seamlessly into background noise. Even when the packet is sent via a public proxy, the callback IP can be redirected elsewhere, separating the trigger path from the control path. Because the router is already positioned at the edge of the network, the moment the reverse shell is established, the operator gains privileged access at one of the most valuable points in the environment.

Campaign evolution: How the operation changed over time

J-magic does not show signs of rapid scaling or infrastructure sprawl. Telemetry from March through September 2024 identified 36 unique IP addresses worldwide that met the signature conditions—less than 0.01% of analyzed NetFlow. This low hit rate reinforces the campaign's priority of precision targeting over mass exploitation.

Victims clustered into two groups:

- **Juniper routers** with self-signed X.509 certificates, indicating that these devices were likely acting as VPN gateways
- **Centrally managed** routers with NETCONF exposed, likely part of larger service provider or telecom environments

Targeting & operational Use

Observed J-magic targets spanned semiconductor manufacturing, energy and solar technology, industrial manufacturing, and IT and network services. Roughly 50% of targeted devices appeared to function as VPN gateways, placing them directly in the authentication and access path for remote users. Others played central roles in routing, filtering, and configuration management—valuable positions for lateral movement or long-term persistence. There is no evidence of destructive activity. The capability appears to be preserved for surveillance, credential access, and future operational leverage rather than for immediate exploitation.

Why this operation matters

J-magic reinforces a critical shift in modern threat operations in which perimeter devices are becoming the payload. By targeting enterprise routers with passive, memory-only malware, attackers gain visibility and access that endpoint defenses never see. There are no alerts, no suspicious processes on user machines, and no obvious indicators until the operator decides to activate the backdoor.

Disruption & defense

Black Lotus Labs detected J-magic by translating malware logic into network-level analytics, identifying the unique packet conditions and correlating them with Juniper device banners to reduce false positives. This detection required upstream visibility into TCP behavior, highlighting why backbone-level observation is essential for uncovering passive, infrastructure-resident threats.

Black Lotus Labs shared IoCs, detection guidance, and hunting recommendations to help organizations identify similar activity in their environments.

Key takeaway

J-magic shows that modern attackers don't always rush the vault. Sometimes they hide nearby—watching traffic, waiting for the signal, and keeping their tools invisible until the moment matters most. When malware lives in routers and activates only on command, the defenders who rely solely on endpoints arrive long after the attack is complete.

Learn more →

Secret Blizzard: How a second crew hijacked the control room and rewrote the job mid-heist.

Over the course of two years, Black Lotus Labs uncovered a longstanding campaign that reads less like a conventional intrusion and more like a takeover.

Beginning in December 2022, a Russian-based threat actor known as "Secret Blizzard" (also referred to as Turla) gained entry to C2 nodes operated by the Pakistan-based actor Storm-0156 (publicly associated with SideCopy and Transparent Tribe). The campaign steadily expanded until November of 2024, when Secret Blizzard successfully infiltrated 33 separate Storm-0156 C2 nodes. As Black Lotus Labs picked up the trail that led from discovering Storm-156's infrastructure, we noticed it was no ordinary operation. Their operators had managed to introduce a physical device into a network at an Indian embassy in Europe and were sending data to a Storm C2.

From that vantage point, Secret Blizzard treated Storm-0156's operations like a pre-built blueprint. They searched what had already been collected, monitored ongoing activity, and selectively inserted their own tooling where the payoff was highest—all while keeping their own infrastructure and fingerprints to a minimum.

The crew: Who's behind the operation

Black Lotus Labs attributes this campaign to Secret Blizzard, a nation-state espionage actor associated with the Russian FSB, and notes the activity aligns with Turla's established pattern of compromising other actors' C2 infrastructure to obscure attribution and harvest intelligence.

Storm-0156 (Pakistan-based) served as the unwitting "first crew," already positioned in regional government targets like Afghanistan and India. Secret Blizzard leveraged that alignment, letting Storm-0156 do the risky access work before swooping in to quietly collect intelligence, exploit trust relationships, and expand operations.

The control room: How the operation was managed

This campaign was managed like a disciplined infrastructure operation. Secret Blizzard maintained access to Storm-0156's C2 estate over long windows and rotated their own C2 nodes as the campaign matured. The Microsoft Threat Intelligence Team (MSTIC) associated Secret Blizzard with a set of VPS IPs used from December 2022 through August 2023 (146.70.158[.]90, 162.213.195[.]129, 146.70.81[.]81), and Black Lotus Labs observed a continuation in May 2024 with C2 rotation to 146.70.158[.]90 and 162.213.195[.]192 lasting through November.

Crucially, Secret Blizzard didn't stop at commandeering servers. They exploited trust relationships to move from Storm-0156 C2s into Storm-0156 operators' workstations (observed in April–May 2023), potentially gaining high-value intelligence like operator tooling, credentials for C2s and victim networks, and archives of prior exfiltration.

The infrastructure stack: how the operation was built

Secret Blizzard's infrastructure in this campaign is best understood as a parasitic multi-layer stack:

- **Storm-0156 C2 layer (host infrastructure):** At least 33 C2 nodes infiltrated, providing visibility into active operations and stored loot.
- **Secret Blizzard overlay C2 layer (VPS nodes):** Used to manage their activity within hijacked infrastructure and reach downstream victims.
- **Downstream victim environments:** Where Secret Blizzard deployed their own malware TwoDash and Statuezy into select Afghan government networks.
- **Operator workstation layer (a higher-value pivot):** Access into Pakistani operator systems enabled appropriation of additional malware families (Waiscot and CrimsonRAT) and possible harvesting of credentials and past collections.

The routes: How the operation moved across the internet

By operating through Storm-0156 infrastructure, Secret Blizzard retrieved files and intelligence without exposing its own toolchain. This tactic could also be used to shift the blame from Secret Blizzard if responders only saw the original actor's infrastructure in incident artifacts.

The campaign's most telling route is the lateral move into operator space. On May 4, 2023, a Pakistan-based dynamic IP (182.188.171[.]52) was observed connecting via RDP to a known Storm-0156 AllaKore C2, while also connecting to Secret Blizzard infrastructure (146.70.158[.]90) during overlapping windows—strongly indicative of compromise of Storm-0156 operators themselves.

Campaign evolution: how the operation changed over time

This operation unfolded in four distinct phases:

- **Phase 1 (Dec 2022–mid-2023):** Initial compromise of a Storm-0156 C2 in December 2022, expansion into additional C2 nodes, and establishment of Secret Blizzard VPS infrastructure linked by MSTIC.
- **Phase 2 (mid 2023):** Downstream deployment into Afghan government networks using TwoDash and Statuezy malware, indicating active exploitation beyond passive collection.
- **Phase 3 (April 2023):** Pivot into Pakistani operator workstations, enabling deeper intelligence collection and tooling theft.
- **Phase 4 (mid 2024):** Expanded use of malware families Waiscot and CrimsonRAT appropriated from Pakistani operator environments; selective engagement with CrimsonRAT C2 nodes previously associated with Indian government/military targeting.

Notably, Black Lotus Labs observed persistent Afghan-network connections again from May to October 2024, with a C2 rotation to 143.198.73[.]108 aligned to prior Storm-0156 infections.

Targeting & operational use

The clearest downstream targeting observed in this campaign focused on Afghan government networks, where Secret Blizzard used Storm-0156's pre-existing access to deploy TwoDash and Statuezy malware and collect data. The campaign also intersected with Indian targeting ecosystems through Secret Blizzard's interaction with select CrimsonRAT C2s, though Black Lotus Labs did not observe Secret Blizzard deploying their own agents into Indian networks. This suggests Secret Blizzard may have limited activity to retrieve data from hijacked C2s or is leveraging existing implants already placed by Storm-0156.

This campaign focused on positioning. Secret Blizzard used Storm-0156's regional footholds to expand intelligence coverage across Afghanistan and potentially adjacent Middle East governmental entities, while extracting value from prior operations already in motion.

Why this operation matters

Secret Blizzard demonstrates that threat actors don't just compromise victims; they also compromise each other. That means "routine" detections (even those that look like cybercrime or a different APT) can be part of a deeper espionage operation.

Disruption & defense

Over the course of this investigation, Black Lotus Labs tracked and documented 37 Secret Blizzard and Storm-0156 C2 nodes, including the malware servers used in their campaign. Our team collaborated closely with MSTIC and used backbone-level visibility to map active and historical control paths. Lumen also blocked traffic across the Lumen global backbone to infrastructure associated with Secret Blizzard and Storm-0156 sub-clusters and ingested campaign indicators into the intelligence feeds supporting Lumen Defender security solutions.

Key takeaway

Secret Blizzard's playbook is a reminder that elite espionage campaigns are often built on stolen staging. They hijack another actor's infrastructure, inherit their access, and quietly extract the intelligence—all while defenders chase the wrong crew. The teams that can see these operations at the infrastructure layer are the ones most likely to stop the attack before it ever breaches the network.

Learn more →

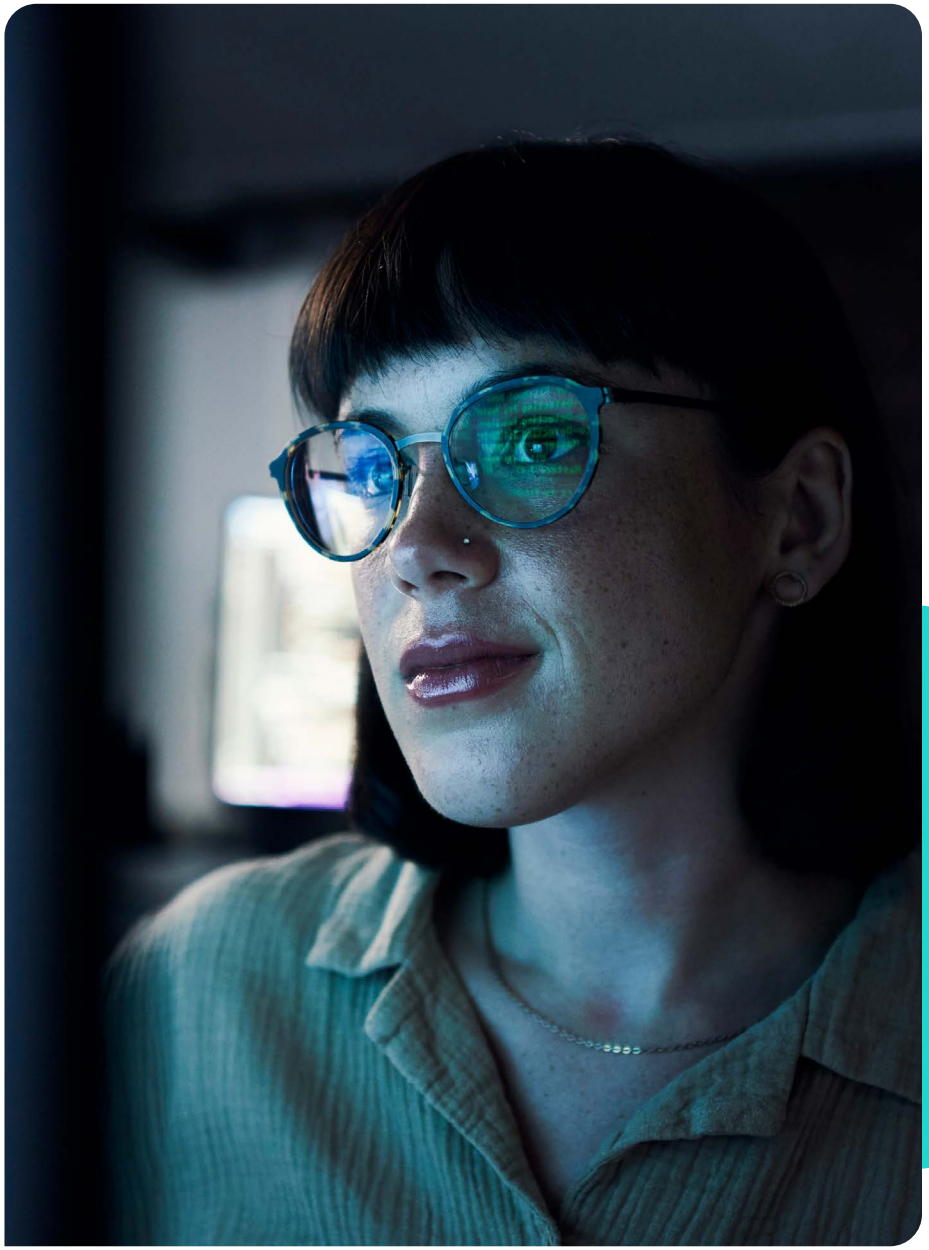

NSOCKS botnet: The proxy crews that turns home routers into disguises

NSOCKS is a criminal proxy business model built on top of a botnet engine, helping threat actors stay anonymous while doing loud things online. Instead of relying on the Tor network or commercial VPNs, NSOCKS gives customers on-demand access to compromised devices that look like real people in real homes.

What makes NSOCKS especially dangerous is its scale and reach. In Lumen telemetry, the service maintains a daily average of 35,000+ bots across 180 countries, with two-thirds of proxies based in the U.S.—a geographic skew that makes it particularly effective for targeting U.S.-based services and organizations.

The crew: Who's behind the operation

NSOCKS presents as a criminal proxy service that has been tied to known threat activity, including use by groups such as Muddled Libra, and is heavily discussed and promoted in criminal communities. Black Lotus Labs links the service to a wider "marketplace" reality where other proxy services (including VN5Socks and Shopsocks5) appear to depend on the same underlying supply chain—meaning one operation can silently power many.

The control room: How the operation was managed

Think of NSOCKS as a storefront with an unusually powerful dispatch layer. Customers don't connect directly to victim devices; they enter through a large layer of "backconnect" C2 nodes that act as the service's entry/exit ramps. Black Lotus Labs assesses over 180 backconnect C2s dedicated to routing and proxying traffic. These nodes signal which bots are available and serve as the customer connection points after purchase.

The service UI also supports targeted selection behavior. Users can filter proxies by domain (including .gov and .edu) and see how many others are currently using a proxy, directly enabling more deliberate targeting choices.

The infrastructure stack: How the operation was built

NSOCKS is primarily powered by the ngioweb botnet. Black Lotus Labs reports that at least 80% of NSOCKS bots in their telemetry originate from ngioweb, which primarily compromises SOHO routers and IoT devices. ngioweb itself is structured like a production pipeline with two major stages:

- **Loader network (initial delivery):** Victims are directed to loader-C2 nodes to retrieve and execute ngioweb malware. Black Lotus Labs couldn't directly observe the initial access vector, but assesses it likely leverages multiple exploits available to the operators. Black Lotus Labs typically tracks 15–20 loader nodes at any given time. Based on supporting research by LevelBlue Labs, Black Lotus Labs suspects access to 10–15 exploits. Loader traffic commonly shows up over port 80 and port 21 (FTP). This arm is assessed to be monitored/controlled by a node at 103.172.92[.]148 which communicates with about half of the tracked loaders.
- **Gatekeeper management layer (DGA C2s):** Once infected, devices reach out to DGA-generated C2 domains that appear to decide whether a bot is worth adding to the proxy pool. Black Lotus Labs typically sees around 15 active domains at a time. These "gatekeepers" connect suitable bots to a backconnect C2, making them available for NSOCKS customers.

The routes: How the operation moved across the internet

NSOCKS's "getaway routes" are its backconnect nodes. When a bot receives a CONNECT command, it reaches out to the supplied backconnect C2 and is instructed to "start proxy," effectively turning the victim into a live relay. The C2 then issues proxy tests—first routing through a "random" URL and later requesting proxy connections to additional domains (including subdomains of nslookups[.]com).

NSOCKS also appears to include a performance/quality mechanism. In some cases, after the proxy starts, a backconnect C2 requests a download of "test.zip"—likely a speed test used to evaluate suitability and potentially pricing/usage value.

Campaign evolution: How the operation changed over time

Black Lotus Labs notes the broader ngioweb botnet has been documented historically dating back to 2018 and 2019, with the tracking and inventory work representing the culmination of over a year's research connecting earlier elements into today's NSOCKS-driven ecosystem.

NSOCKS itself first appeared in the fall of 2022. According to Spur, it likely operated under LuxSocks prior to that. Based on older ngioweb domains/files, Black Lotus Labs suspects LuxSocks was also largely powered by ngioweb.

Targeting & operational use

NSOCKS is designed for abuse by default. Customers can purchase IPs with cryptocurrency for 24 hours and use them for fraud, reconnaissance, spam, and phishing-related activity.

The victim pool also tells an important story. About 40% of bots remain active for a month or longer, providing a stable window for repeated misuse. Black Lotus Labs also observed co-habitation, where serious groups (such as the APT group Pawn Storm) have been found abusing the same devices as ngioweb, meaning infected routers can become shared infrastructure across multiple threat actors simultaneously.

Why this operation matters

NSOCKS is a prime example of the industrialization of "anonymity-as-a-service." Not only did NSOCKS have 35,000 malware-backed proxies, its ecosystem is also engineered to let many different threat actors run their own operations through the same machinery—even enabling the creation or feeding of other proxy brands/services.

It also highlights a defender's trap. Many NSOCKS endpoints appear as ordinary residential IPs, allowing threat actors to blend into traffic patterns that often bypass simplistic geo/ASN heuristics. Additionally, NSOCKS' ability to filter for sensitive domains like .gov and .edu makes "precision abuse" easier.

Disruption & defense

Using global internet visibility, Black Lotus Labs traced both active and historical C2 nodes, including previously undiscovered infrastructure that has been in use since mid-2022. Lumen then blocked all traffic across its global network to/from the dedicated infrastructure associated with the ngioweb botnet, and is releasing IoCs to help others detect and disrupt the operation.

Key takeaway

NSOCKS is the cybercrime equivalent of a crew that never has to show its face. It rents disguises at scale, is built on a botnet supply chain (ngioweb), and leverages a global backconnect "switchboard" to route traffic through unsuspecting homes. And because the service's real advantage is infrastructure (reach, stability, and reuse), defenders have to go beyond endpoints—using upstream visibility to identify the routes attackers rely on to blend in.

Learn more →

Raptor Train botnet: The nation-state botnet with an enterprise-grade command center

If most botnets feel like smash-and-grab crews, Raptor Train is a long-running operation with a plan, platform, and pipeline. Black Lotus Labs first discovered the large, multi-tier botnet of SOHO and IoT devices in mid-2023 while investigating compromised routers. We believe that it was operated by nation-state Chinese threat actors known as Flax Typhoon.

Over four years in the making at the time of its takedown, Raptor Train was designed with an enterprise-grade control system that gave its operators the ability to deploy its full range of capabilities from one integrated panel. This control panel allowed a greater number of users with less specialization to conduct effective campaigns.

At its peak in June 2023, Raptor Train consisted of over 60,000 actively compromised devices. Over the course of Raptor Train's operation, Black Lotus Labs observed more than 200,000 compromised SOHO routers, NVR/DVRs, network attached storage (NAS) servers, and IP cameras pulled into the botnet, making it one of the largest Chinese state-sponsored IoT botnets discovered to date. One of its C2 domains was so heavily queried that it even cracked both Cisco Umbrella and Cloudflare Radar "Top 1 Million" popularity lists—a concerning milestone since "popular" domains can sometimes slip past defenses through whitelisting behavior.

The crew: Who's behind the operation

Black Lotus Labs used multiple converging signals to link Raptor Train to Flax Typhoon: management and operational timeframes were consistent with Chinese working hours (Monday–Friday); targeting aligned to Chinese strategic interests; research showed Chinese language use in tooling; and there were overlapping Tactics, Techniques, and Procedures (TTPs) between Raptor Train and other Chinese nation-state groups.

The operational picture suggests a nation-state program built for scale and optionality, with an infrastructure that supports reconnaissance, exploitation, and potentially future disruptive actions.

The control room: How the operation was managed

Raptor Train's standout feature is its robust, enterprise-grade controller the actors call "Sparrow." The operators run the network through a series of distributed payload and C2 servers, a centralized Node.js backend, and a cross-platform Electron front-end capable of managing upwards of 60 C2 servers and their infected nodes at any time.

Sparrow supports an entire suite of operator workflows, including the scalable exploitation of bots, vulnerability and exploit management, remote management of C2 infrastructure, file upload/download, remote command execution, and the ability to tailor IoT-based DDoS attacks at scale. It also enables automation and steady log/bot data collection to increase operators' situational awareness. This frees Raptor Train's crew up for hands-on exploitation while the platform itself handles orchestration.

Black Lotus Labs also identified an additional Tier 3 Sparrow management node called "Condor," designed to support exploitation workflows (payload generation, exploit attempts, verification, and logging) and to assist in discovery and testing activities (including 0-days).

The infrastructure stack: How the operation was built

Black Lotus Labs observed at least three tiers of activity:

- **Tier 1: The "crew" (compromised SOHO/IoT devices)**

Raptor Train's Tier 1 footprint spans modems/routers, IP cameras, NVR/DVR devices, and NAS systems. Black Lotus Labs assessed operators likely exploited 20+ device types, using a mix of 0-day and n-day (known) vulnerabilities.

Black Lotus Labs has dubbed the primary implant on Tier 1 nodes "Nosedive." It is a custom Mirai variation compiled for major SOHO/IoT architectures (MIPS, ARM, SuperH, PowerPC, etc.). Nosedive was typically deployed from Tier 2 payload servers through a unique URL encoding scheme and domain injection method that encodes the requested C2 domain and combines it with a "key" identifying the bot and its architecture. Nosedive droppers then inject that key into the Nosedive implant payload that is deployed to the Tier 1 node. Once deployed, Nosedive ran in-memory only and allowed the operators to execute commands, upload and download files, and run DDoS attacks on compromised devices.

Critically, Nosedive and associated droppers are memory-resident only and deleted from disk. Black Lotus Labs observed additional anti-forensics techniques, including running-process-name obfuscation, multi-stage infection chains, and killing remote management processes—making detection and forensics that much harder. The lack of persistence in most Nosedive implants is also telling. Operators appeared unconcerned with long-term persistence because the pool of vulnerable devices is so large that they can simply rotate bots continuously to meet operational needs.

- **Tier 2: The "equipment caches" (exploitation servers, payload servers, and C2 servers)**

Tier 2 consists of dedicated virtual infrastructure used to deliver payloads and coordinate infected devices. Payload servers fall into first-stage (longer-running, generic retrieval) and second-stage (often hosted on high, random ephemeral ports like 32123/38525) systems. Second-stage servers typically appear in more targeted efforts, possibly to better obfuscate sensitive exploitation such as 0-days.

From 2020 to 2022, there were approximately 1–5 C2 nodes. That number jumped to 11 in mid-2023, 30 from February to March 2024, and over 60 from June to August 2024. Each time we identified a growth in C2 nodes, we observed an increase in Tier 1 nodes (bots).

- **Tier 3: The "control layer" (management nodes, sparrow nodes, and broader network architecture)**

Tier 3 nodes manage Tier 2 manually via SSH (port 22) and, for Tier 2 C2 nodes, automatically via TLS (port 34125). Manual management sessions were observed almost exclusively during Chinese working hours (Monday–Friday), while the Sparrow-driven 34125 connections were consistent across all hours and days, reflecting persistent automation.

The primary Sparrow interface is named "节点综合控制工具v1.0.7" which translates to "Node Comprehensive Control Tool v1.0.7." It includes permission-based management that supports a team of operators executing commands, moving files, collecting data, and managing DDoS tasking across infected devices and distributed infrastructure.

In short, operational tasks originate from Tier 3 Sparrow management nodes, route through Tier 2 C2s, and are delivered to Tier 1 bots. This pyramid-style structure creates a huge, disposable base with increasingly specialized control layers above it.

The lifecycle differences are part of the design. Tier 1 bots last an average of 17 days, consistent with physical device churn and power cycling. Tier 2 and Tier 3 nodes are typically procured VPS infrastructure with longer lifespans averaging 77 days. Tier 2 VPS nodes are geographically distributed worldwide, while Tier 3 servers are largely based in Hong Kong or the PRC.

The routes: How the operation moved across the internet

Raptor Train's "routes" are its tiered command pathways. Sparrow management nodes issue tasks that flow through the Tier 2 C2 layer and into tens of thousands of rotating Tier 1 bots. Payload delivery relies on the unique encoding and injection method to bind bots to specific C2 infrastructure, while Tier 2/Tier 3 VPS longevity provides a stable backbone that survives as Tier 1 devices churn. This design (disposable endpoints combined with durable control) complicates network defense. Removing one layer or device isn't enough, as the botnet spans tens of thousands of devices.

Campaign evolution: How the operation changed over time

Black Lotus Labs divided Raptor Train into four campaigns, each showing evolution in scale, naming patterns, infrastructure rotation, and targeting emphasis:

- **Crossbill (May 2020–April 2022):** Began with root domain k3121.com, later expanding to encoded random alphanumeric subdomains.
- **Finch (July 2022–June 2023):** Signified by root domain b2047.com, first registered in September 2019. Subdomain patterns evolved in June 2023 to longer strings. By mid-June 2023, at least 10,000 distinct devices were infected.
- **Canary (May–August 2023):** More tailored targeting (notably ActionTec PK5000 modems, Hikvision IP cameras, Shenzhen TVT NVRs, ASUS RT-*/GT-* routers). The Canary campaign marks the first observed use of multi-stage droppers and some in-memory "persistence," and a major scale jump: Tier 2 C2 growth (from 1–3 active to 10+) and Tier 1 growth (from 10,000 to 60,000+).
- **Oriole (June 2023–August 2024):** Signified by root domain w8510.com, with broad device expansion observed April–August 2024 (VNPT iGate routers; AXIS IP cameras; and compromised NAS devices such as QNAP NAS, Zyxel NAS, Fujitsu NAS and Synology NAS.) By August 2024, Raptor Train maintained an average of 30,000 Tier 1 compromised devices despite the 17-day churn rate. w8510.com entered Cisco Umbrella rankings in June 2024 and Cloudflare Radar's Top 1 Million by August 2024.

Targeting & operational use

Black Lotus Labs uncovered targeting activity concentrated on U.S. and Taiwanese entities in military, government, higher education, telecommunications, defense industrial base (DIB), and IT sectors. In late December 2023, operators conducted extensive scanning targeting the U.S. military, U.S. government, IT providers, and DIBs. The team also observed widespread global targeting, including a government agency in Kazakhstan, along with more targeted scanning and likely exploitation attempts against vulnerable technologies like Atlassian Confluence and Ivanti Connect Secure appliances (likely via CVE-2024-21887) in the same strategic sectors.

While Black Lotus Labs has not observed DDoS attacks originating from Raptor Train, the botnet clearly maintains DDoS capability, and Black Lotus Labs assesses this is a capability the operators may preserve for future use.

Why this operation matters

Raptor Train is a clear signal of how nation-state operations are evolving away from single-purpose intrusions and toward reusable infrastructure programs. Its architecture is built for resilience through churn (Tier 1 rotates fast), stability through VPS control layers (Tier 2/3 persist longer), and scale through professional orchestration (Sparrow and Condor). The inclusion of a C2 domain in Cisco Umbrella and Cloudflare Radar "Top 1 Million" also underscores how infrastructure scale can create defensive blind spots, especially when "popular" domains are treated as inherently less suspicious.

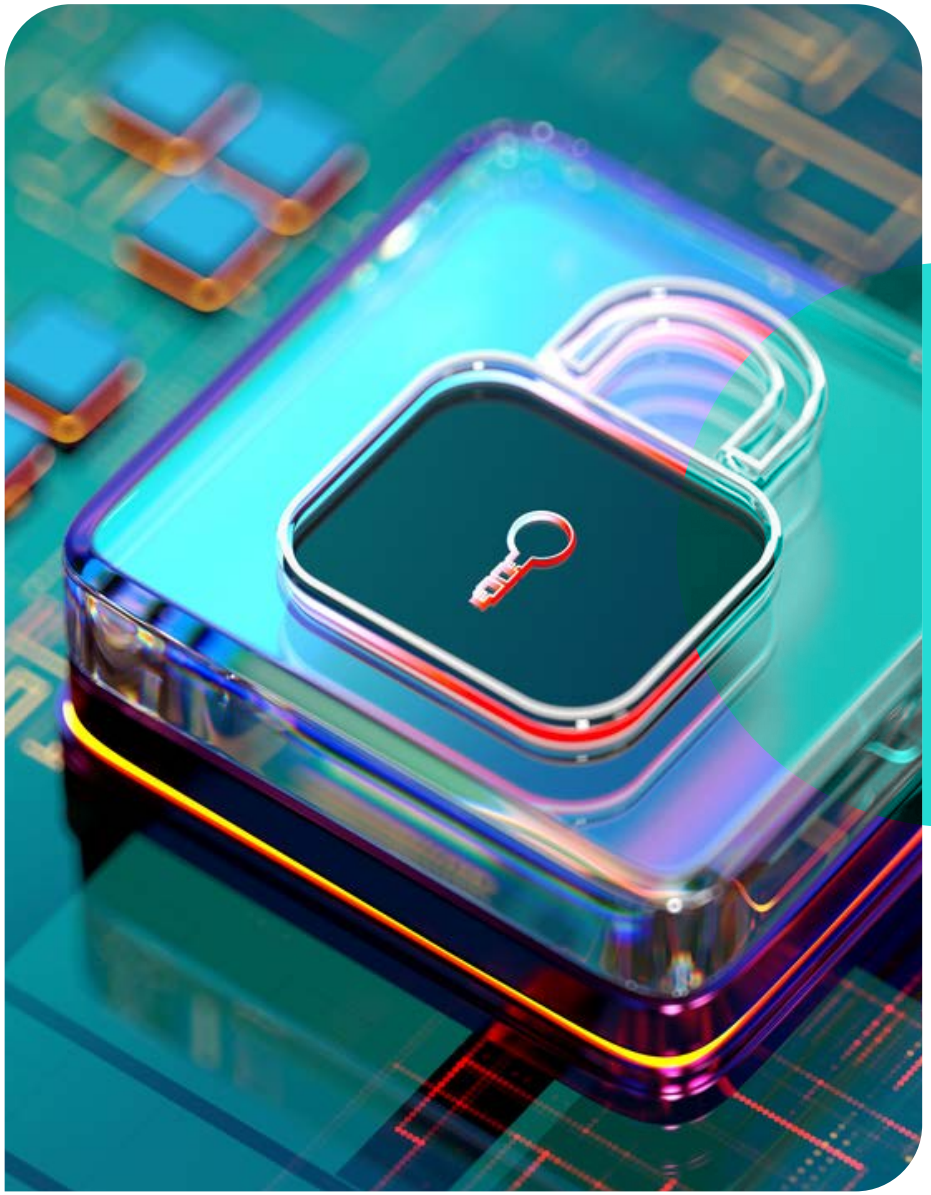

Disruption & defense

Black Lotus Labs shared threat intelligence to warn agencies across the U.S. Government about emerging risks to strategic assets. Lumen also null-routed traffic to known infrastructure used by Raptor Train operators, including distributed botnet management, C2, payload, and exploitation infrastructure.

Key takeaway

Raptor Train was more than a botnet. It was a nation-state logistics system built over years, scaled across hundreds of thousands of devices, and coordinated through an enterprise-grade control plane. Raptor Train demonstrates what modern campaigns look like when the infrastructure layer becomes the operation, and why defenders need network intelligence to proactively spot, and stop, attacks.

Learn more ⟶

LUMEN®

2026 threat predictions

In 2026, the most dangerous cyber operations won't look radically different at the moment of impact. Breaches will still begin with stolen credentials, exploited edge devices, or trusted infrastructure abused at scale. However, what will change is how fast those operations are assembled through the proliferation of generative AI tooling, how little forensic evidence they leave behind, and how effectively attackers blend into the background noise of the internet itself.

From Black Lotus Labs's vantage point inside a global internet backbone, we believe four shifts are already coming into focus.

Prediction 1: Setup gets faster as adoption of generative AI and agents goes mainstream

Speed is a crucial element of every good heist. In 2026, threat actors won't just move fast at execution time—they'll plan and assemble infrastructure at machine speed.

Across nearly every campaign in this report—J-magic, Secret Blizzard, SystemBC, DanaBot, and multiple proxy botnets—the network edge was the preferred point of leverage. We expect this trend to continue in 2026 with a sharp acceleration in chained exploit paths targeting edge devices and internet-exposed management interfaces, including firewalls, VPN gateways, routers, and orchestration panels. Crucially, we also expect attackers to pair these chained exploit paths with agentic AI systems that automate lateral movement decisions in real time.

Instead of manually pivoting from one compromised system to the next, we believe that attackers will increasingly rely on AI-driven agents that evaluate privilege levels, identify adjacent trust relationships, select the next best exploit path, and adapt tactics mid-operation based on network response. For defenders, this means they won't be facing a single exploit or malware family—they'll be confronting living attack chains that reconfigure as conditions change.

Edge devices are already prized because they sit at the crossroads of authentication, encryption, and routing and often lack deep forensic visibility. When agentic AI accelerates how quickly attackers move through those devices, the window for detection shrinks dramatically. Defenders will need to be prepared to counter this shift with more proactive visibility.

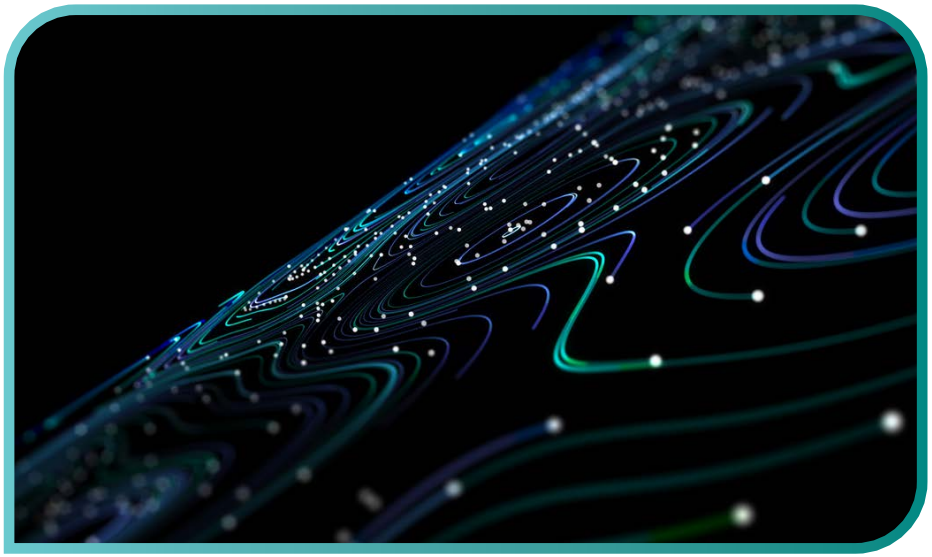

Prediction 2: Targeting will focus on opportunity at the edge

Given the success with which adversaries targeted edge devices and services in 2026, we believe that attackers will continue to shift toward opportunistic exploitation of internet-exposed infrastructure in 2026.

Lumen backbone telemetry shows that malicious scanning for edge devices and exposed services is continuous and indiscriminate, with continuous malicious scanning representing 33% of all malicious traffic blocked by Lumen Defender. Firewalls, VPN gateways, remote management interfaces, identity services, and load balancers are probed relentlessly across all sectors. When a reachable device is found—especially one with weak authentication, missing patches, or limited logging—attackers move quickly, using whatever technique is most effective at the moment. This trend highlights the importance of monitoring and securing vulnerable edge devices and services.

EoL devices will remain a preferred target given the near-endless supply of devices that are aging out of manufacturer support. These systems often lack modern security controls and occupy privileged positions in the network—making them durable footholds once compromised. As long as they remain exposed, they will continue to attract attention from both criminal and nation-state actors.

For defenders, this means risk in 2026 will be defined by exposure. Organizations that prioritize asset visibility, patch discipline, privilege management, and retiring unsupported devices will be better positioned to anticipate attacks.

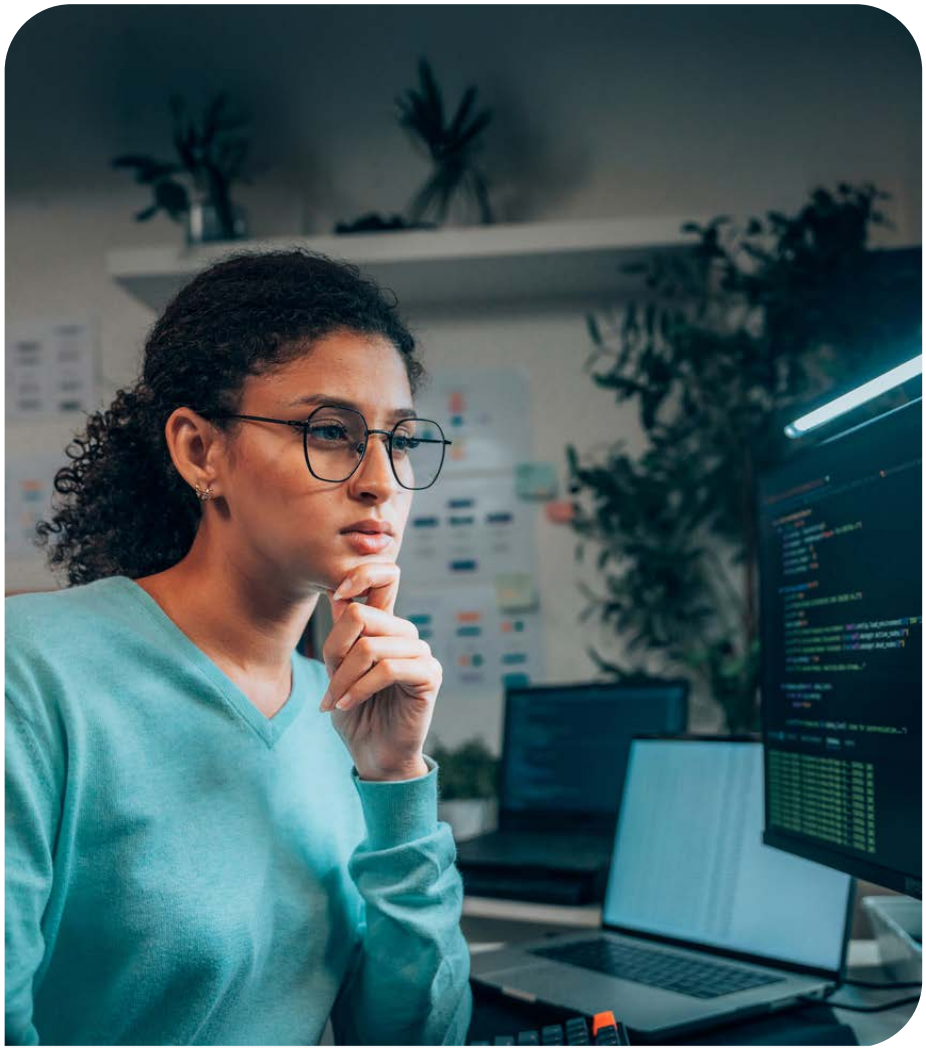

Prediction 3: The real signals lie in the network

By the time defenders investigate a compromised firewall or router, the most important evidence will already be gone. That's why we believe detecting adversary networks in addition to individual tools will be essential in 2026. This is because attackers are increasingly targeting devices with limited forensic capabilities, including limited logging, short memory, in-memory or ephemeral malware, privileged access, and high throughput.

From Lumen's perspective, the earliest signals won't come from device telemetry. They'll emerge from how infrastructure behaves collectively, such as:

- Rapid rotation of C2 nodes
- Sudden emergence of proxy layers
- Traffic patterns that reveal orchestration rather than usage
- Coordination across geographies that no single enterprise can see

This is where backbone-level visibility becomes decisive. The ability to see relationships will be essential to spotting attacks upstream. Threat actors are already designing campaigns that assume defenders will investigate endpoints too late, as seen in the Threats Deconstructed section above. In 2026, the defenders who win will be the ones who enhance and complement EDR and perimeter defense efforts with network infrastructure intelligence to spot attacks while they're still in the planning and preparation stage.

Prediction 4: The best disguises will be legitimate infrastructure

The most effective heists don't involve dramatic escapes. Instead, a well-prepared crew simply walks out the front door unnoticed.

In 2026, we predict that adversaries will rely even more heavily on malware-backed proxy networks, SOHO-based botnets, hijacked VPS infrastructure, and "clean-looking" residential IP spaces to hide their operations in plain sight. As seen in Black Lotus labs threat research, these tactics obscure attribution; blur the boundaries between criminal and nation-state activity; enable shared infrastructure across multiple campaigns; and allow attackers to rent, trade, or reuse capabilities at scale.

The result is a threat landscape where the attacker infrastructure itself becomes the capability. When cybercriminals blend into legitimate traffic flows, traditional indicators fail, reputation-based blocking lags behind reality, and ASN and geolocation filters lose relevance. To uncover rapidly emerging threat actor capabilities and stop crews in their tracks, defenders must have tools in place to spot the subtle signs of a new attack infrastructure taking shape.

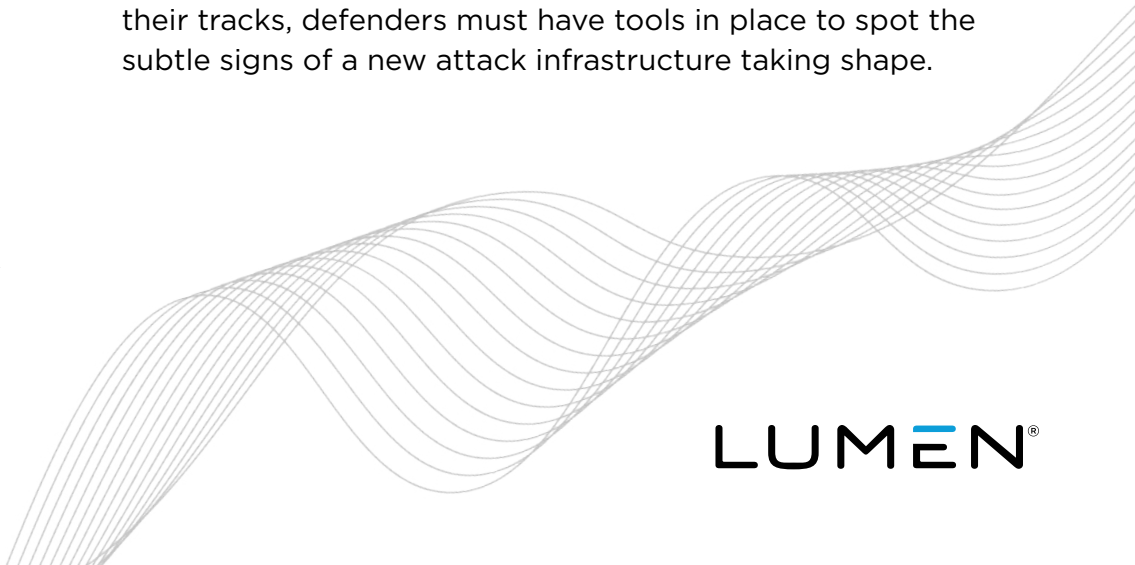

LUMEN®

Black Lotus Labs data illustrates how today's attackers are using malware-backed proxy networks to route traffic from any area that's proximal to their end target.

Blocked attack IPs by country

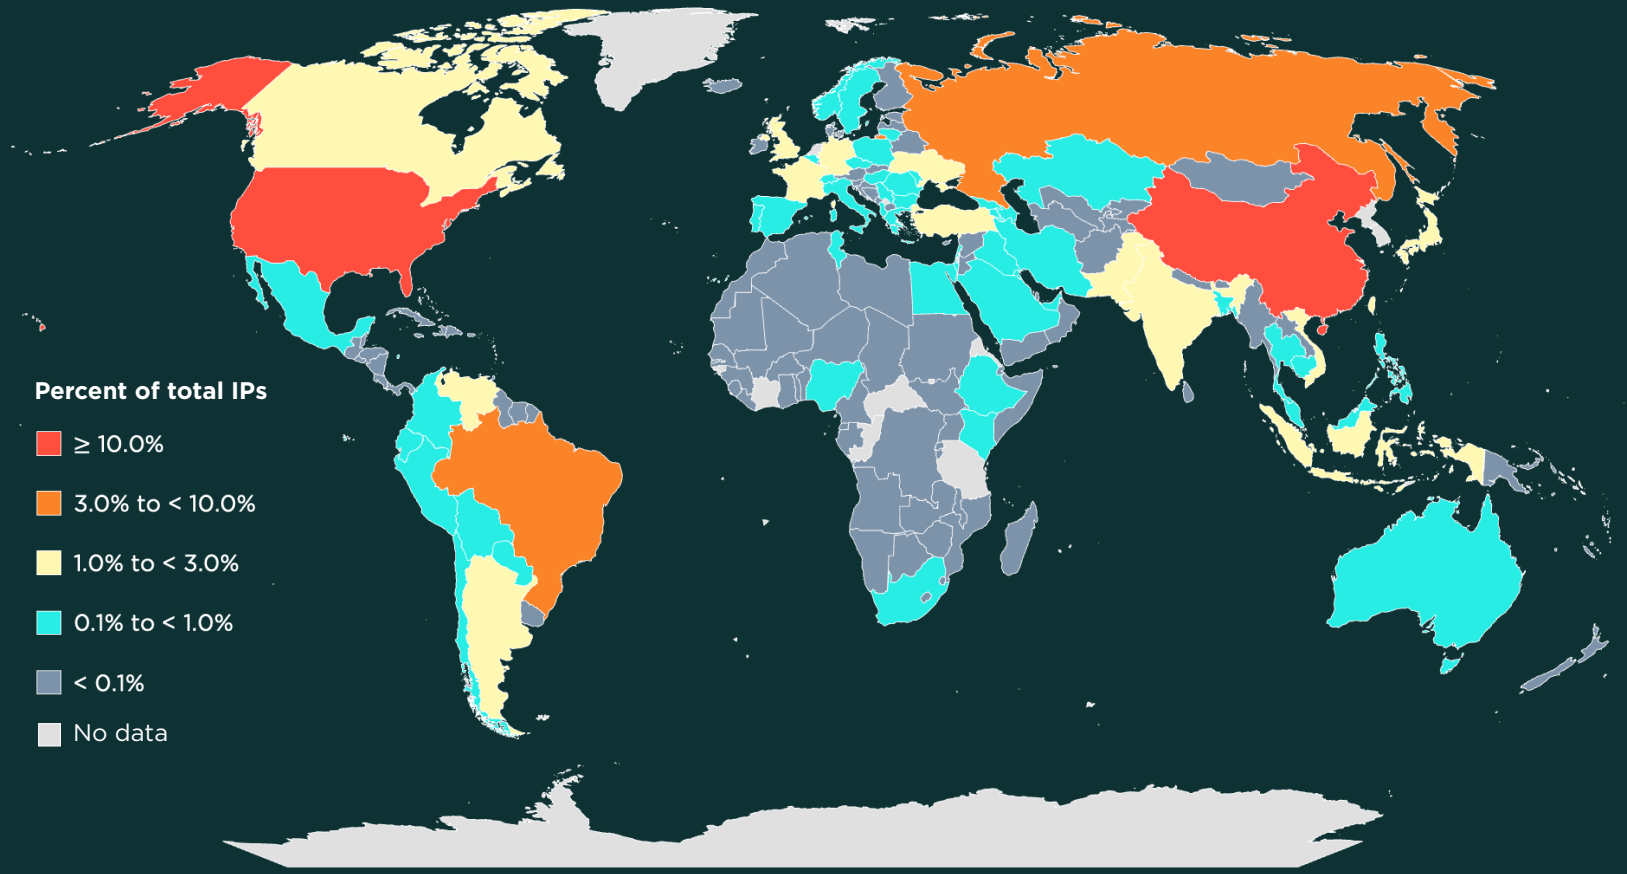

This map shows an elevated level of attacker activity from IPs located in North America, illustrating how threat actors leverage proxy IPs located close to their targets.

LUMEN®

Geographic distribution of C2 IPs

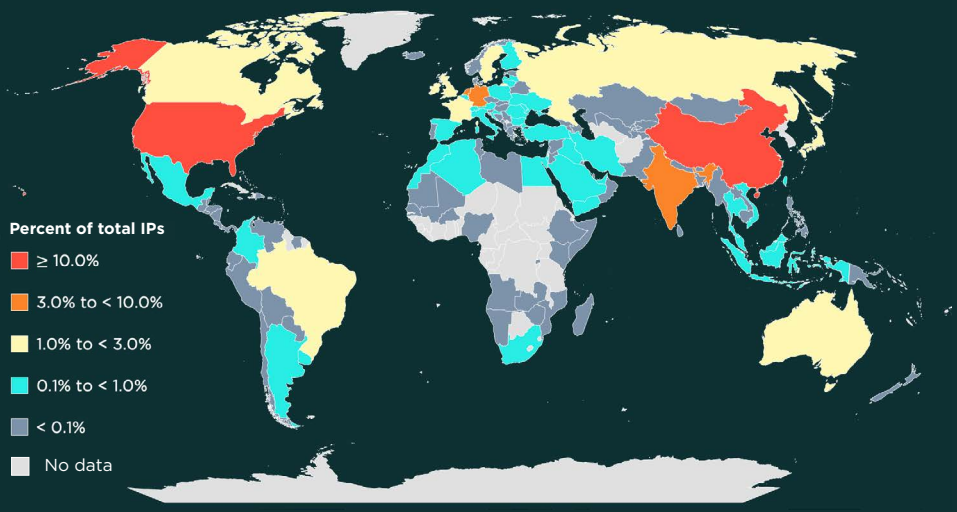

This map highlights where command and control servers are located, including elevated levels of activity in Western Europe and North America. This indicates a preference toward placing administrative infrastructure in locations less likely to be flagged as malicious, especially when compromising corporate networks.

Geographic distribution of bot and victim IPs

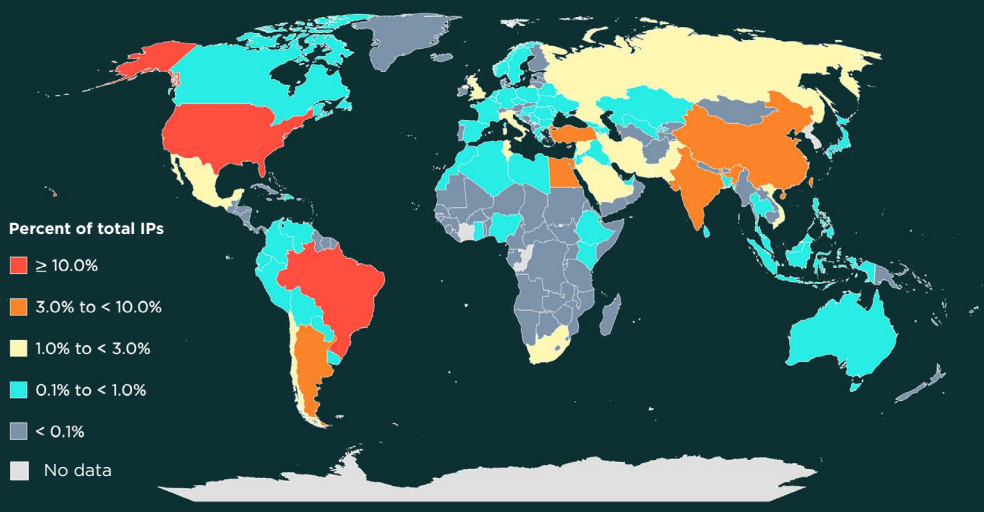

This map shows the locations of all compromised IPs detected by Lumen. The elevated concentration of compromised IPs in South America and the Middle East illustrates the global reach of malicious botnets.

Defense guidance: Stopping the heist before it happens

Modern threat operations are no longer defined by a single exploit, piece of malware, or intrusion event. They are built as systems that are assembled over long periods of time, tested in pieces, and activated only when conditions are favorable. As a result, effective defense in 2026 will depend less on detecting the final payload and more on disrupting the infrastructure, routes, and preparation phases that make large-scale operations possible.

The guidance below reflects patterns observed across Black Lotus Labs's telemetry—from nation-state espionage campaigns to criminal proxy ecosystems—and highlights where defenders can most effectively shift the balance.

1 Defend the edge like it's the vault door

Across nearly every campaign in this report—J-magic, Secret Blizzard, SystemBC, DanaBot, and multiple proxy botnets—the network edge is the preferred point of leverage. Firewalls, VPN gateways, routers, and management interfaces offer three things attackers value most: privileged access, long uptime, and limited forensic visibility.

For defenders, this means that they need to treat gateway edge devices as high-value assets, not infrastructure plumbing. Security teams should apply the same rigor to patching, monitoring, and access control that they would to domain controllers or crown-jewel servers. Teams should also assume that attackers will increasingly chain multiple edge weaknesses together—initial access, lateral movement, and persistence—before any endpoint alert fires.

Practical actions:

- ☐ Thoroughly inventory internet-exposed services and management interfaces, including shadow IT and legacy devices.
- ☐ Monitor for anomalous authentication attempts and configuration changes on edge devices, even when traffic originates from residential or "benign" IP space.
- ☐ Plan for detection techniques that do not rely on host-based agents, especially for appliances and network gear.

2 Shift from indicators to infrastructure awareness

IoCs still matter, and Black Lotus Labs frequently publishes new IoCs on our GitHub. However, IoCs arrive late in the attack lifecycle—oftentimes after threat actors have already reached the perimeter. Black Lotus Labs's research repeatedly shows that infrastructure behavior like C2 relationships, proxy usage, traffic patterns, and routing dependencies reveal malicious activity far earlier than malware signatures. This is especially critical as attackers rotate infrastructure rapidly, reuse shared proxy ecosystems, and operate on devices with limited or no local logging.

For defenders, this means that detection has to focus on relationships rather than just IPs or hashes. If 2025 has taught us anything, it's that seeing how systems connect can be more valuable than knowing what tool is running. After all, adversaries can change tools overnight. However, it's a lot harder to hide the infrastructure they rely on to operate at scale.

Practical actions:

- ☐ Look for unusual data flows, including large outbound transfers to nearby geographic regions or unexpected cloud providers.
- ☐ Track how traffic enters and exits the environment, not just whether it is "known bad."
- ☐ Correlate authentication activity, proxy usage, and outbound connections to identify emerging campaigns before tooling is fully deployed.

3 Treat proxy networks as active threat infrastructure

One of the clearest trends across criminal and nation-state activity is the industrialization of proxy services. Malware-backed proxy networks built from SOHO devices, IoT systems, and compromised VPS infrastructure allow attackers to blend into normal traffic and bypass traditional geofencing or reputation-based controls.

As shown in campaigns like NSOCKS, 5socks, and SystemBC, these networks go beyond anonymity layers—acting as operational platforms that support phishing, brute force, data exfiltration, and ransomware. For defenders, this means that they can no longer afford to treat the residential IP space as a trust signal. A "clean" IP reputation today does not mean safe tomorrow.

Practical actions:

- ☐ Monitor for suspicious activity originating from residential and VPS IP ranges, especially against authentication services.
- ☐ Actively identify and block open proxies and known malicious proxy services where possible.
- ☐ Recognize that attackers may deliberately accept higher detection rates during reconnaissance phases, then switch to cleaner proxy infrastructure for exploitation.

4 Assume blurred lines between crime and espionage

Secret Blizzard illustrates a critical reality of modern threats: the same infrastructure can support both cybercrime and nation-state espionage, sometimes simultaneously. Threat actors are increasingly co-opting each other's access, tools, and data—making attribution less important than impact.

For defenders, this means that they need to treat all unauthorized access incidents as potentially strategic.

Practical actions:

- ☐ Investigate compromises with the assumption that access may be resold, reused, or repurposed.
- ☐ Monitor for secondary activity after initial containment, especially lateral movement and data staging.
- ☐ Prioritize visibility into east-west traffic and unusual operator-driven behavior inside the network.

5 Use scale against the attacker

Attackers succeed by operating at scale. Defenders can regain the advantage by doing the same.

Black Lotus Labs's ability to detect, correlate, and disrupt activity across global NetFlow demonstrates how upstream visibility changes the equation—allowing defenders to observe campaigns forming, not just detonating.

For defenders, this means that they need to expand the scope of their visibility beyond individual assets to the network paths connecting them.

Practical actions:

- ☐ Leverage upstream telemetry, intelligence sharing, and automated response where possible.
- ☐ Integrate network-level detections with identity, endpoint, and cloud signals.
- ☐ Reduce attacker dwell time by blocking malicious infrastructure early, even when confidence is still emerging.

Conclusion

The central lesson of 2025—and the defining risk of 2026—is this: modern cyber threats are built long before they are launched. Adversaries are rapidly assembling infrastructure ecosystems powered by generative AI, resilient proxy networks, compromised edge devices, and shared criminal-state platforms. By the time an alert triggers inside the enterprise, the groundwork for the attack has already been laid.

The advantage now belongs to defenders who can see beyond individual malware samples or isolated intrusion events and instead recognize coordinated infrastructure behavior as it forms. Exposure defines risk. Edge devices, identity systems, and unmanaged or EoL assets have become the preferred leverage points, meaning that detection has to evolve from reactive investigation to proactive disruption at the network layer.

In 2026 and beyond, the decisive battleground will not be the endpoint alone—it is the malicious infrastructure that makes modern campaigns possible. Organizations that combine strong cybersecurity fundamentals with upstream visibility and infrastructure intelligence will be best positioned to shrink attacker dwell time, block malicious systems early, and stop adversarial operations before they ever reach the vault.

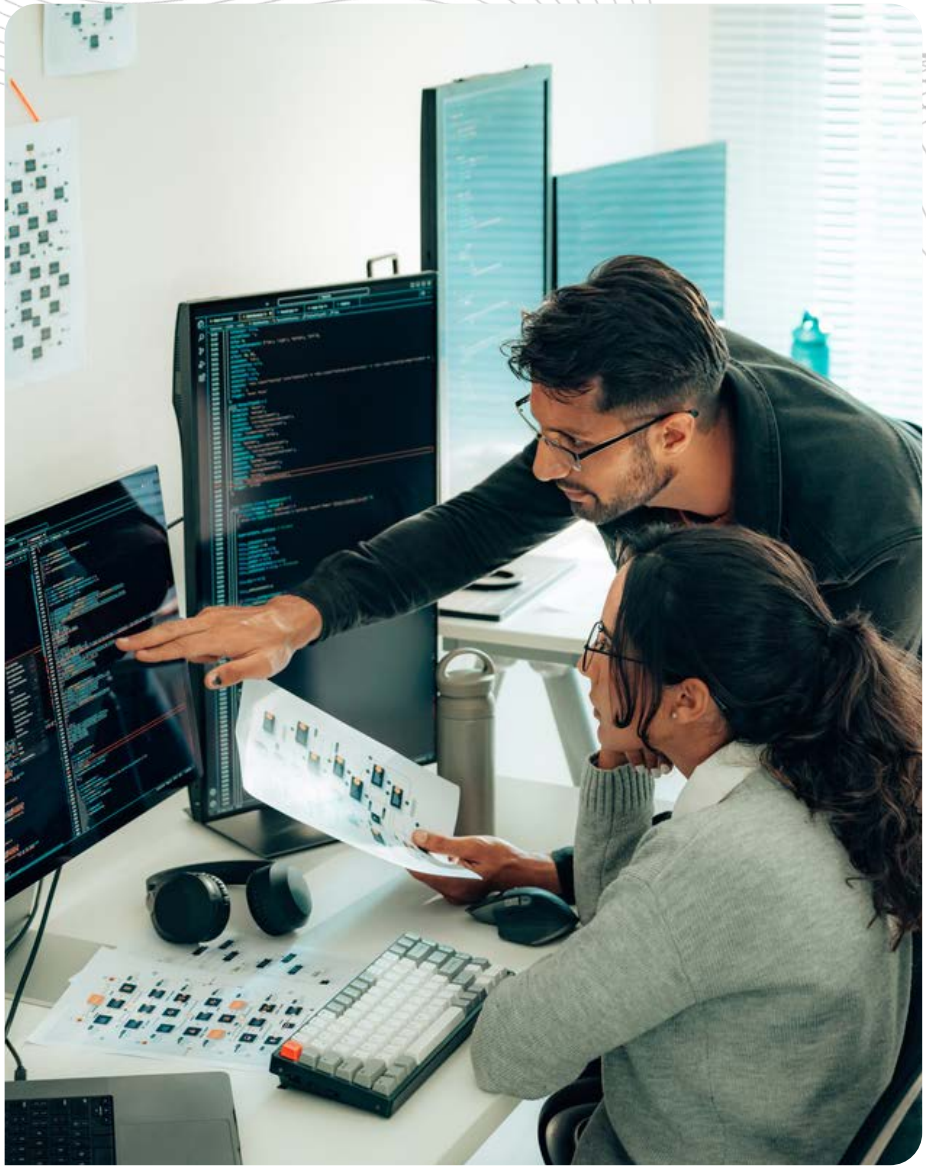

LUMEN®

Inside Lumen global internet backbone visibility

Understanding how attackers build, evolve, and conceal their operational ecosystems across the global internet is essential to anticipating their movements and neutralizing their impact before threats materialize inside the enterprise.

Whereas many other security tools detect damage after it hits a device, Black Lotus Labs sees the assembly of the tools used to cause it—observing adversaries as they choose hosts, test persistence, spin up new servers, and abandon old ones. We then use this insight to disrupt attackers upstream, null-routing malicious infrastructure before it can ever be weaponized.

As the threat research and operations arm of Lumen, Black Lotus Labs operates from a vantage point few organizations in the world possess—offering direct, continuous visibility into an interconnected global internet backbone. Our transit traffic provides visibility to 99% of all public IPv4 addresses, allowing Black Lotus Labs to observe how adversary infrastructure is built, tested, reused, and abandoned—often before a single victim is publicly identified.

Why backbone visibility changes the threat equation

Black Lotus Labs' unique visibility into a vast portion of global internet traffic allows us to see malicious infrastructure as it forms, not after it strikes. This unlocks earlier detection of botnets, C2 systems, and nation-state activity that remain invisible to traditional endpoint and perimeter-based defenses.

Black Lotus Labs is designed to detect the unique signal adversaries leave across the internet as they traverse from their networks to the networks of their victims through multiple hops around the world.

That visibility has fueled a sustained record of industry-shaping discovery and disruption for over a decade. Black Lotus Labs has uncovered and disrupted critical threat ecosystems spanning large-scale DDoS operations and ransomware-as-a-service infrastructure, introducing the world to the first evidence of PRC nation-state use of home routers to disguise espionage as a remote worker tunneling across the internet, going on to reveal dozens of campaigns.

Cited by key government leaders, Black Lotus Labs partnered with the US government and private industry to discover Volt Typhoon's intrusion into US-based critical infrastructure, discovered the first ever Volt Typhoon 0-day, and has disrupted major PRC campaigns from KV-botnet to Raptor Train, seeding the key visibility for law enforcement action.

This upstream perspective is especially critical in 2026, as attackers increasingly spin up new infrastructure at AI-powered machine speed, operate on internet-exposed edge devices with limited forensic visibility, rely on residential and VPS-based proxy networks, blur the line between criminal and nation-state infrastructure, and rapidly rotate tools to evade detection.

Turning visibility into defense—without chasing alerts

Of course, seeing the heist is only useful if it changes the outcome. The same intelligence that powers Black Lotus Labs's research also feeds operational defenses across the Lumen network. This is where products like Lumen Defender Essentials & Plus, Lumen DDoS Mitigation, Lumen Defender Threat Feed, Lumen Defender Advanced Managed Detection & Response (AMDR), and many more Black Lotus Labs-powered solutions enter the story.

Beyond discovery, Black Lotus Labs helps actively degrade adversary capability at internet scale. Leveraging its position on the global backbone, our team disrupts malicious infrastructure through techniques such as null routing—simply put, cutting off attacker servers so they can no longer communicate with infected devices across the global backbone. These actions, often executed in coordination with law enforcement, have dismantled long-running criminal proxy networks and significantly impaired major botnets, translating intelligence into real-world impact. By combining upstream visibility, decisive disruption, and partnership with authorities, Black Lotus Labs helps keep the internet cleaner and measurably raises the cost of operating malicious infrastructure worldwide.

However, our customers don't have to wait for global, partnered disruptions to take down a criminal network or nation-state espionage campaign. Black Lotus Labs feeds this advanced intelligence into Lumen Defender security solutions, so as soon as Black Lotus Labs can detect, our customers can get protection.

This model reflects the same lesson repeated throughout this report. The earlier defenders disrupt infrastructure, the less damage an adversary can do. Lumen approach does not replace endpoint, identity, or cloud security. It complements them by delivering additional intelligence so defenders can spot and block attacks before they reach the point of intrusion.

Research methodology

This report is informed by a comprehensive review and synthesis of Black Lotus Labs research published between September 2024 and January 2026. The analysis draws on original investigations, global network telemetry, and campaign-level findings documented across Black Lotus Labs content during that period. Findings were further supplemented with original research to identify key 2025 trends and criminal patterns.

Black Lotus Labs, Lumen threat research division, produces threat intelligence to help protect businesses and maintain internet security. This intelligence is integrated into Lumen Defender security solutions. Using data from the Lumen network and advanced machine learning, Black Lotus Labs identifies malicious hosts and infrastructure. Each host is given a Risk Score reflecting its severity. Our goal is to be accurate and timely enough to protect customers, reduce false positives to spare customer analysis and resources, and use the threat intelligence product for additional research.

The factors influencing Risk Scores, and the methods by which we determine the threats that are expressed as metrics in this report, are explained here:

How is a risk determined?

The Lumen Defender Risk Score is calculated through an aggregate of three primary components:

- **Severity**
- **Confidence**
- **Temporal Decay**

Severity

Severity is quantified via multiple sub-factors:

- **Threat category:** Determined by adversary progression along the "Kill Chain." Early-stage threats, such as "Scanning," may represent minimal or serious risk, depending on the correlation to the individual source. Conversely, advanced-stage threats, such as high-volume (by byte count) communications with adversarial C2 servers, may indicate data exfiltration and represent significant risk.
- **Threat progression:** Categories include Proxy, Scan, Phish, Malware, Bot, Attack, and C2, mapped to their respective stages in the Kill Chain.
- **Asset correlation:** Reputation reporting at the IP Address level must consider the underlying network architecture. For single-server IP addresses, event-to-IP correlation is high. For hosting/CDN IP addresses, correlation weakens due to multiple assets sharing the address.

For instance, while comparing June–August metrics to those in September–December of 2025, UDP/ICMP-based hostile activity fell by about 83%, with its share of overall hostile traffic dropping from 19.8% to 5.0%. Here, our fidelity is good enough to show micro trends, in this case one that points to a marked reduction in classic reflection and amplification behavior, and a shift towards more targeted, primarily TCP-based attacks on specific services and applications.

Confidence

Confidence is determined by:

- **Source confidence:** Each intelligence source is evaluated and assigned a confidence factor based on historical accuracy and reputation. High-confidence sources contribute more weight to the overall risk score.

  As an example, during the June–December 2025 window, our systems recorded approximately 4.63 billion High, Very High and Severe risk "events." Of these, roughly 3.65 billion events—78.7%—are backed by attack or scan confidence scores of at least 90. As stated above, "events" are considered as any traffic sent to a customer that was malicious enough to be blocked by Defender. As these events carry strong detection quality, they also represent a low likelihood of false positives in high priority alerts.

- **Entity validation:** Upon reporting a new malicious entity, the Threat Research Team attempts validation through sandboxed surrogate assets and reverse engineering. Responsive entities increase the confidence factor.

  One example of this is the creation and use of our "validators," which simulate an active victim to analyze and map its behavior.

Time-based risk score decay

Risk scores are dynamically reduced as threat indicators age or cease to exhibit malicious behavior. Due to the transient nature of internet-based threats and frequent host relocation by adversaries, risk scores for IP addresses are adjusted downward over time to reflect current threat posture.

As we see activity fall under the risk threshold, for instance, normal traffic to null-routed IP addresses is reinstated.

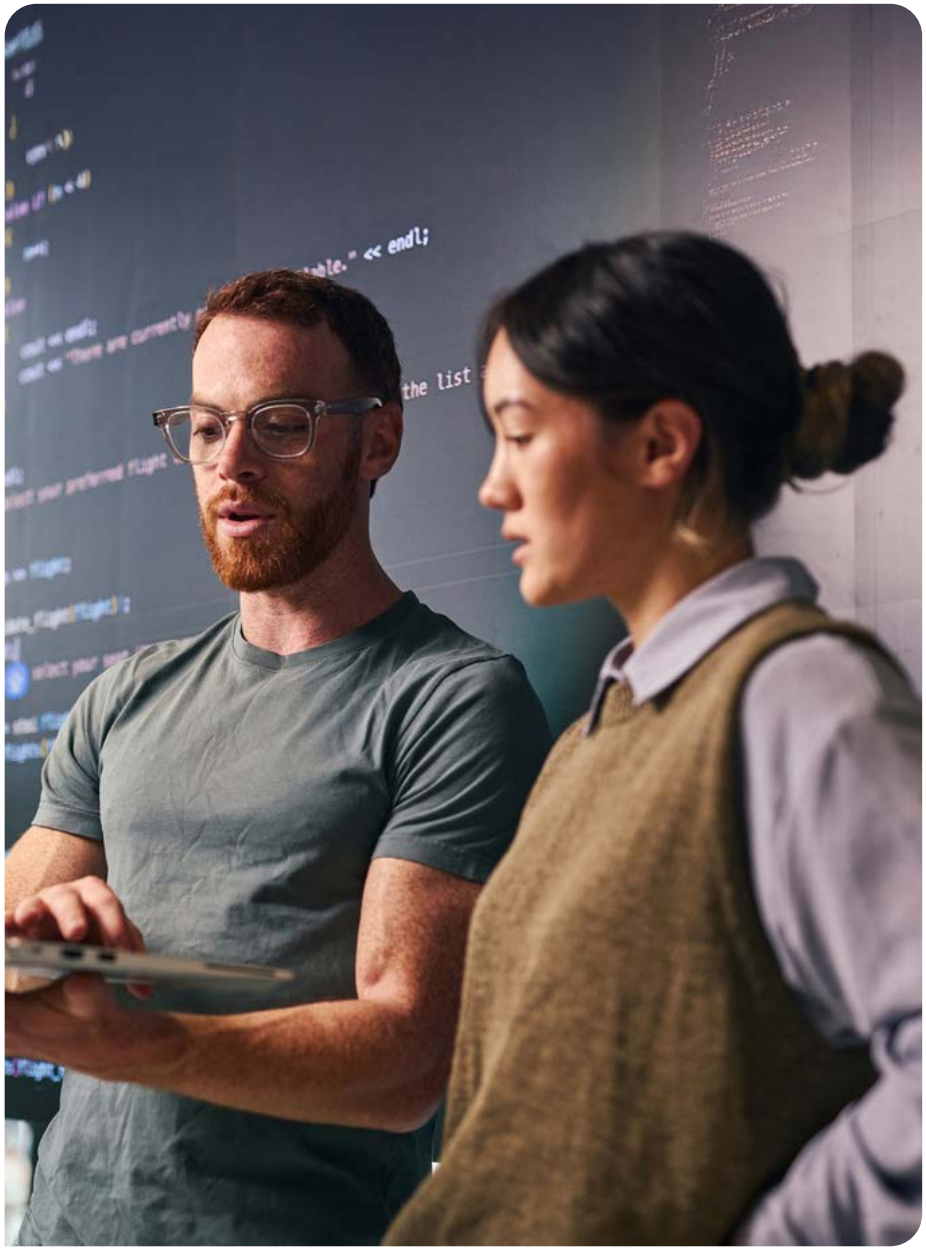

LUMEN®

BLACK LOTUS LABS® by Lumen

Footnotes

1. The Center for Applied Internet Data Analysis (CAIDA), AS Rank, January 2025.
2. "Navigating the Paths of Risk: The State of Exposure Management in 2024," XM Cyber, Cyentia Institute (2024).
3. "Countering Chinese State-Sponsored Actors Compromise of Networks Worldwide to Feed Global Espionage System," Cybersecurity & Infrastructure Security Agency (September 3, 2025); "Cyber agencies unveil new guidelines to secure edge devices from increasing threat," National Cyber Security Centre (February 4, 2025); "Mitigation strategies for edge devices: Practitioner guidance," Australian Signals Directorate (February 4, 2025); "Security considerations for edge devices (ITSM.80.101)," Government of Canada (February 4, 2025); "Security considerations for edge devices," Communications Security Establishment Canada, Canadian Centre for Cyber Security (2025); "Securing Network Infrastructure Devices," Cybersecurity & Infrastructure Security Agency (September 6, 2006).
4. Barrero, Jose Maria, Nicholas Bloom, and Steven J. Davis, 2021. "Why working from home will stick," National Bureau of Economic Research Working Paper 28731.
5. Tushar Subhra Dutta "Danabot Malware Resurfaced with Version 669 Following Operation Endgame," Cyber Security News (November 12, 2025).

866-352-0291 | lumen.com | info@lumen.com