

LUMEN®


BLACK LOTUS LABS®
by Lumen

Healthcare **report**

The 2026 Lumen Defender Threatscape Report

Why visibility at breach misses the plot



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 healthcare organizations—where patient safety, clinical continuity, and regulatory compliance are paramount—it's now critical to gain visibility into these operations before they disrupt care delivery systems or expose sensitive patient data.

Modern cyber operations look less like isolated break-ins and more like carefully staged heists. Long before ransomware is deployed or a breach impacts patient records, 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 healthcare organizations only detect these activities once they intersect with perimeter defenses or clinical systems, by which point the operational groundwork has already been laid.

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

What are edge devices?

Edge devices—such as routers, switches, firewalls, and 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.

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.**

This shift has significant implications for the healthcare industry. While endpoint detection and response (EDR) and traditional security controls remain essential, attackers are increasingly bypassing these defenses by exploiting exposed edge infrastructure—particularly in environments with distributed care delivery, internet of medical things (IoMT) devices, third-party integrations, and hybrid cloud architectures. These threats then trickle down, compromising patient privacy and safety. In a sector where uptime directly impacts patient outcomes and regulatory obligations are constantly evolving, managing exposure at the edge is now as critical as protecting core systems.

Informed by insights from Black Lotus Labs, this healthcare report breaks down the key trends we observed in 2025, outlines our top predictions for 2026, and provides actionable steps to help healthcare security leaders disrupt adversarial activity earlier—before it impacts clinical operations, patient safety, or compliance posture.

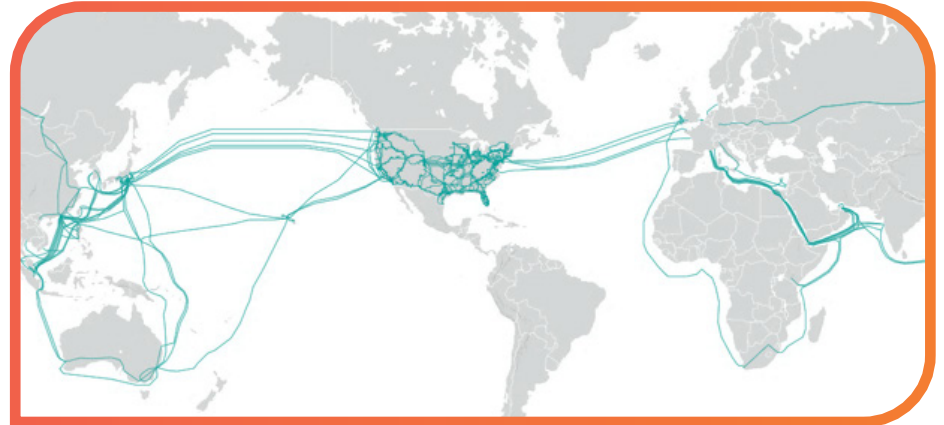
[Download the complete 2026 Lumen Defender Threatscape Report for additional insight](#) →

1. The Center for Applied Internet Data Analysis (CAIDA), AS Rank, January 2025.

Inside Lumen's internet backbone

Lumen manages and operates one of the largest, most connected, and most deeply peered networks in the world.¹ Black Lotus Labs provides:

- 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
- 5,000+ C2s disrupted in 2025



[View live threat analytics](#) →

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, small office/home office (SOHO) devices, Internet of Things (IoT) and IoMT hardware, and virtual servers blending malicious activity into everyday traffic. For healthcare organizations, this constant churn complicated attribution and overwhelmed traditional detection models—especially across environments with distributed care delivery, third-party connectivity, and legacy systems.

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 end-of-life (EoL) devices to create moving targets. SOHO and IoT/IoMT devices allowed attackers to “live off the land” while disguising operations behind residential IP space and trusted network pathways. Shared use of compromised infrastructure further blurred the lines between criminal and state activity, increasing the risk of missed detections and delayed response—particularly in healthcare environments where uptime pressures can delay investigation or containment.

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

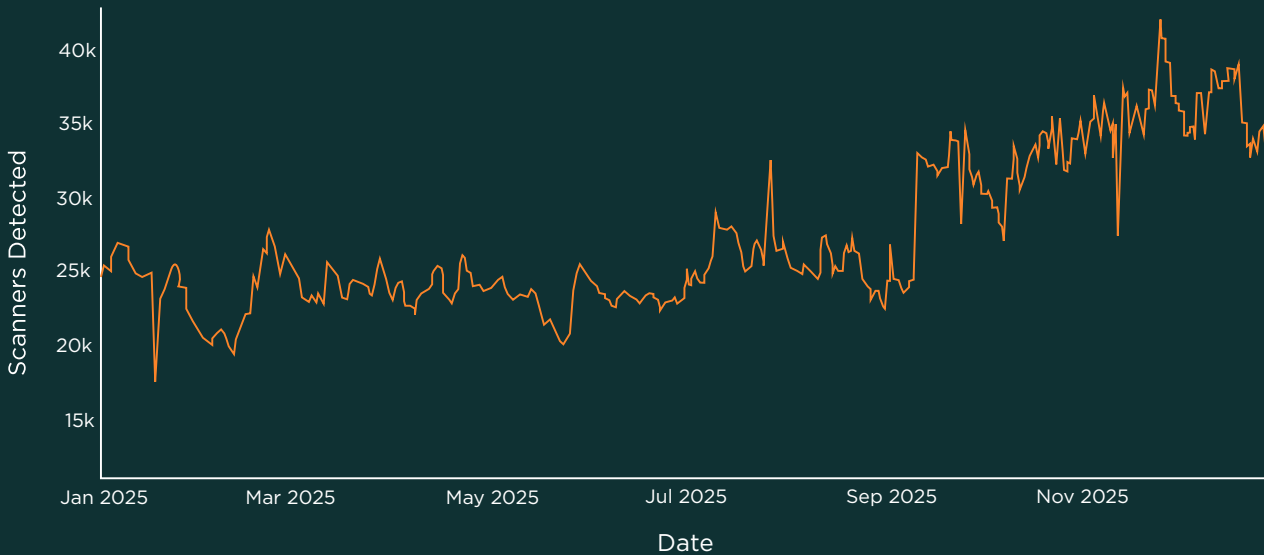
1 **Generative AI changed the tempo, unleashing attacks at machine speed.**

Threat actors embraced automated tasking, scanning, infrastructure rotation, and C2 management in 2025, in addition to adopting new generative AI-enabled attack vectors. This automation sustains malicious campaigns indefinitely with minimal oversight while continuously adapting to new targets.

For healthcare organizations, this acceleration is particularly concerning. AI is being used to generate highly convincing phishing campaigns, automate credential harvesting, and rapidly identify exposed systems across clinical and administrative environments. Human operators increasingly step in only to refine targets—often focusing on high-impact systems such as electronic health records (EHRs), remote access tools, or third-party vendor connections.

This scale and speed forced healthcare security leaders to rethink how they detect and disrupt attacks—shifting from reactive response to identifying early-stage signals before ransomware, data exfiltration, or operational disruption occurs.

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 and Plus between October–December 2025.

2 Attackers moved 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, firewalls, and connected medical devices that underpin healthcare delivery and interoperability.

Here, attackers can “live in the middle,” harvesting credentials, monitoring traffic, and quietly expanding access across interconnected systems. In healthcare environments—where networks span hospitals, clinics, labs, imaging systems, and remote care platforms—this creates opportunities to move laterally without triggering traditional controls.

Sophisticated attackers know that completing the heist in a single sitting may alert defenses. By lurking in devices outside the reach of standard security controls and reaching back out days, weeks, or even months after the initial access, attackers can delay action—timing their operations to maximize disruption, often during critical care windows. This makes it significantly harder for defenders to correlate activity and respond before clinical systems or patient data are impacted.

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. We saw adversaries assemble polished platforms—complete with customer support and subscription tiers—as a way to scale access, automate abuse, and lower the barrier to entry for cybercrime.

This professionalization lowers the barrier to entry for targeting healthcare organizations, enabling less-skilled actors to launch ransomware campaigns, credential harvesting operations, and supply chain attacks against providers, payors, and life sciences companies.

For healthcare CISOs, this means threats are no longer limited to highly skilled adversaries. Instead, scalable, service-based attack platforms can rapidly target organizations that rely on complex ecosystems of vendors, partners, and interconnected systems—amplifying risk across the entire care delivery network.

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/IoMT devices, and high-volume virtual private server (VPS) hosts became rentable identities—on-demand personas that can be purchased for a few dollars' worth of cryptocurrency. This tactic allowed attackers to disguise themselves as trusted remote employees, bypassing geofencing, ASN-based blocking, IP reputation checks, and many Zero Trust location signals.

This tactic is especially problematic in healthcare, where remote access, telehealth, and third-party integrations are essential. Attackers can impersonate trusted users, vendors, or even clinical staff—bypassing geofencing, IP reputation checks, and other access controls designed to protect sensitive systems.

As a result, traditional identity and access safeguards are increasingly undermined, raising the risk of unauthorized access to patient data, clinical systems, and operational platforms.



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. Attribution became less about ownership and more about intent.

For healthcare organizations, this convergence raises the stakes. The same infrastructure used for ransomware or data theft today may also support espionage, intellectual property theft, or disruption of critical healthcare services tomorrow.

This dynamic is particularly relevant for life sciences organizations and large healthcare systems, where sensitive research, patient data, and operational systems represent high-value targets. As attribution becomes less clear, defenders must focus less on who is behind an attack and more on identifying and disrupting the infrastructure enabling it.

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.

For healthcare organizations, this upstream visibility is critical. It enables security teams to identify and disrupt malicious infrastructure before it is used to target patient care systems, expose protected health information (PHI), or disrupt operations—shifting defense earlier in the attack lifecycle.

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, which is now powered by generative AI.

This shift is especially urgent for healthcare security leaders. In an environment defined by connected devices, distributed care delivery, regulatory pressure, and patient safety risks, defending infrastructure is now foundational to maintaining clinical continuity and trust.

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 the 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, we expect to see a sharp acceleration in AI-enabled chained exploit paths targeting edge devices and internet-exposed management interfaces, including routers, VPN gateways, firewalls, identity systems, and the connected technologies that support clinical operations, remote care, and third-party integrations. We believe 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.

Edge devices are already prized because they sit at the crossroads of authentication, encryption, and routing and often lack deep forensic visibility. Agentic AI simply accelerates how quickly attackers can move through those devices.

For healthcare organizations, this compresses the window between initial exposure and material impact—whether that's ransomware deployment, disruption of clinical systems, or unauthorized access to PHI. As attacker speed increases, proactive visibility becomes critical to maintaining operational resilience and protecting patient care.



Prediction 2: Targeting will focus on opportunity at the edge

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 Essentials & Plus. 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.

For healthcare organizations, this means risk will increasingly be defined by exposure. Environments with large numbers of connected devices, legacy systems, and third-party integrations (common across hospitals, payors, and life sciences organizations) expand the attack surface and create more opportunities for exploitation. Organizations that prioritize asset visibility, patch discipline, identity and access management, and the retirement or isolation of unsupported systems will be better positioned to reduce risk.

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.

From our perspective, the earliest signals won't come from device telemetry. They'll emerge from how infrastructure behaves collectively, such as rapid C2 rotation, sudden proxy layer emergence, orchestration traffic patterns, and coordination across geographies.

For healthcare organizations, augmenting endpoint and perimeter defenses with network-level intelligence will be critical. This enables earlier detection of malicious campaigns before they disrupt clinical operations, impact patient care, or lead to regulatory and compliance incidents. In complex healthcare environments where data flows across providers, payors, partners, and devices, this broader visibility is essential to understanding and mitigating risk.



Prediction 4: The best disguises will be legitimate infrastructure

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. 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. This increases the risk of unauthorized access to sensitive systems and patient data, often without immediate detection.


For healthcare organizations, this creates a particularly difficult challenge. Remote access, telehealth services, vendor connectivity, and distributed care delivery all rely on legitimate external traffic—making it harder to distinguish malicious activity from normal operations. In this environment, detecting the subtle emergence of malicious infrastructure before it’s used at scale will become a critical defensive capability.



Defense guidance: Stopping the heist before it happens

Modern threat operations aren't 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.

The guidance below reflects patterns observed across Black Lotus Labs's telemetry and highlights where defenders can most effectively shift the balance.



1 Defend the edge like it's the vault door

Firewalls, VPN gateways, routers, and management interfaces offer privileged access, long uptime, and limited forensic visibility. Defenders should treat these gateway edge devices as high-value assets and assume that attackers will chain multiple edge weaknesses together—initial access, lateral movement, and persistence—before any endpoint alert fires.

- 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

Infrastructure behavior like C2 relationships, proxy usage, traffic patterns, and routing dependencies reveals malicious activity far earlier than malware signatures. Detection should focus on relationships rather than just IPs or hashes.

Priority 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

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. Security teams can no longer treat the residential IP space as a trust signal.

Priority 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

2025 showed us that the same infrastructure can support both cybercrime and nation-state espionage. Threat actors are increasingly co-opting each other's access, tools, and data—making attribution less important than impact. Defenders must treat all unauthorized access incidents as potentially strategic.

Priority 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. Security teams can regain the advantage by doing the same. Organizations must expand the scope of their visibility beyond individual assets to the network paths connecting them.

Priority 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.



The central lesson of 2025—and the defining risk for healthcare organizations in 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 and IoT devices, and shared criminal-state platforms. Organizations that combine strong cybersecurity fundamentals with upstream visibility and infrastructure intelligence will be best positioned to shrink attacker dwell time, disrupt malicious activity early, and prevent attacks before they impact clinical operations, patient safety, or sensitive health data.

Lumen combines global internet traffic visibility with original research through Black Lotus Labs to spot malicious infrastructure as it forms—unlocking earlier detection of botnets, C2 systems, and nation-state activity. This same intelligence feeds operational defenses across the Lumen network, so as soon as Black Lotus Labs can detect, our customers can get protection.

Our approach does not replace endpoint, identity, or cloud security. It complements them by delivering additional intelligence so healthcare organizations can better detect and disrupt threats earlier in the attack lifecycle. The result is a more proactive security posture that helps protect patient data, maintain clinical continuity, and support resilience in an increasingly complex, interconnected, and highly regulated environment.

