LUMEN®

BLACK LOTUS LABS® by Lumen

The 2026 Lumen Defender Threatscape Report

CISO Update

Established cyber threats are becoming faster, broader, and more persistent with AI



Table of contents

Executive summary 03

Mid-year scorecard 05

Constant pressure, sharp bursts, and a shorter path to impact 06

The decisive signal may appear before the endpoint alert 08

Attack infrastructure is becoming regenerative 10

H2 2026 defense posture checklist 13

Accelerating defense to the speed of attack 14

Recommendations for Lumen customers 15

Executive summary

Lumen's global internet backbone gives Black Lotus Labs®, our threat research team, a unique view into how malicious infrastructure navigates the public internet. The expert research is paired with machine learning and automation to identify and disrupt emerging threat infrastructure before it reaches organizations. With this vantage point, we publish an annual Threatscape Report—and this mid-year update—to translate our network-level intelligence into practical insights to help security leaders navigate threats that grow faster, broader, and more persistent.

When Lumen released its *2026 Threatscape Report* earlier this year, our network-level perspective informed four key predictions:

1. Setup of adversarial infrastructure will accelerate as adoption of generative AI and agents becomes mainstream.
2. Attackers will focus on opportunistically exploiting internet-exposed infrastructure, including edge devices and services.
3. As attackers increasingly target devices with limited forensic capabilities, defenders will need to augment existing detection and response tooling and perimeter defense efforts with network infrastructure intelligence to proactively spot and block attacks upstream.
4. Adversaries will continue leveraging malware-backed proxy networks, Small Office/Home Office (SOHO) botnets, hijacked Virtual Private Server (VPS) infrastructure, and "clean-looking" residential IP spaces to obscure operations, impersonate legitimate users, and bypass geolocation and Zero Trust controls.

While this forecast largely held true, the greatest changes we observed were in attackers' speed, scale, and staying power.

AI and automation are compressing the attack sequence—from vulnerability discovery to scanning, exploitation, bot enrollment, and monetization. Mandiant's *M-Trends 2026 report* found that the mean time to exploit vulnerabilities is now negative, dropping to an estimated -7 days.[1] Black Lotus Labs has observed similar pre-disclosure exploitation at the network layer—including ransomware activity targeting a **CheckPoint vulnerability** three weeks before public disclosure and the **FortiBleed campaign**, where automated brute-force and exploit activity shifted rapidly across attacker infrastructure.

FortiBleed is particularly instructive for security leaders because it demonstrates the challenge of defending against machine-speed campaigns. Over three weeks, Black Lotus Labs observed an always-on operation that targeted more than **10,000 Fortinet endpoints** and generated more than **65 million attacks**. Given the persistent flow of attacks, even a temporary misconfiguration could lead to compromise. As AI and automation help reduce the time, labor, and expertise required to move from discovery to exploitation, reactive security operations become increasingly difficult to scale.

1. Jurgen Kutscher, **"M-Trends 2026: Data, Insights, and Strategies From the Frontlines,"** Google Cloud Blog (March 23, 2026)

The first half of 2026 reinforced this dynamic across the threat landscape. Threat actors leverage persistent scanning and edge-device exploitation to search for reachable opportunities continuously. Residential, SOHO, and VPS infrastructure give attackers readily available places to hide, rotate, and rebuild. Large-scale distributed-denial-of-service (DDoS) botnets demonstrate how quickly adversarial capacity can shift under pressure. Crucially, much of this activity develops before a traditional detection-and-response workflow ever sees an active threat inside the enterprise.



Bottom line for CISOs?

Rather than asking the SOC to react faster to every alert, the focus should instead be on shifting defensive action earlier in the attack lifecycle. Organizations should complement detection and response with continuous posture and exposure management, behavior-based network threat intelligence, and policy-driven automation. Network intelligence adds a preventative layer by identifying how threat actors behave across the internet, creating opportunities to block their infrastructure upstream even while AI-assisted threats accelerate.

Mid-year scorecard

The operating model anticipated in The 2026 Lumen Defender Threatscape Report is becoming more entrenched, requiring organizations to respond with greater creativity, proactivity, and coordination across their security layers.

1 Setup gets faster ↑

Automated brute force, bursts of scanning and exploit activity against specific edge services, and rapid infrastructure rotation continue to shorten the path from discovery to impact. AI may accelerate research, targeting, and orchestration, but it is an emerging layer rather than proof of a wholly new threat vector.

CISO implication: Automate exposure reduction and response wherever attackers are already operating at machine speed.

2 Targeting follows opportunity at the edge ✓

Scanning against Fortinet FortiGate Firewalls, Cisco ASA, Palo Alto GlobalProtect, and other exposed devices remained steady, with sharp spikes around emerging vulnerabilities. Attackers are more focused on whether an exploitable device or service is reachable than the industry an organization operates in.

CISO implication: Prioritize internet-facing firewalls, VPNs, routers, and management interfaces immediately after crown-jewel systems, with equivalent discipline around patching, access control, exposure reduction, and monitoring.

3 The real signals lie in the network ✓

Network backbone visibility is a valuable source of proactive threat intelligence, enabling organizations to identify early warning signs like targeted communication patterns, infrastructure clusters, and changes in attacker behavior before a vendor disclosure or endpoint alert.

CISO implication: Pair existing detection and response tools and perimeter defenses with network infrastructure intelligence to detect attacker infrastructure as campaigns form and block malicious traffic upstream.

4 The best disguises are legitimate infrastructure ✓

Residential proxies, SOHO botnets, compromised VPS systems, and shared IP pools continue to make malicious activity look geographically and reputationally normal. Several proxy networks tracked by Black Lotus Labs rebuilt after disruption.

CISO implication: Treat clean reputation scores, familiar geography, and residential sources as context rather than proof of trust.

Faster

Less time between vulnerability discovery, scanning, exploitation, and monetization.

Black Lotus Labs observed dedicated exploit servers targeting Ivanti Connect Secure services **within 24 hours** of disclosure (CVE-2026-10520).

Within **72 hours**, more than **100** dedicated servers began targeting Ivanti services.

Broader

Continuous internet-wide pressure, punctuated by bursts around disclosures and zero-days.

Black Lotus Labs detected an average of nearly **4,000 brute force attacks** and exploit scanners per day from January–August 2026.

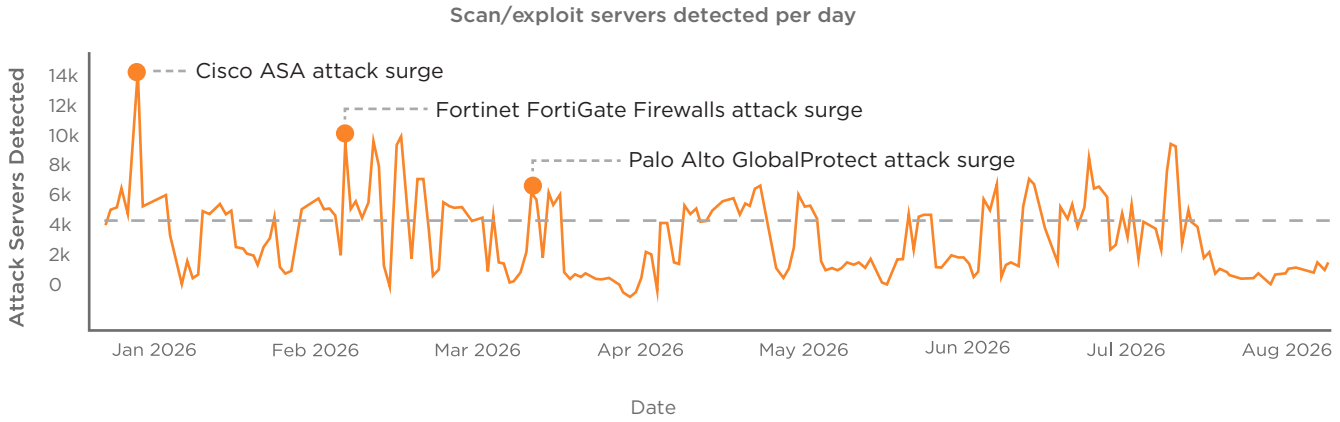
More persistent

Bot, proxy, and control infrastructure is replenished and reconstituted after disruption.

The IPIDEA botnet reconstituted nearly **112%** of its ecosystem less than 6 months after large-scale disruption.

Constant pressure, sharp bursts, and a shorter path to impact

Adversaries are continuously inventorying internet exposure. In 2026, we've seen a steady baseline of malicious scanning and brute force activity, punctuated by short, opportunistic surges around newly disclosed vulnerabilities and zero-days. That combination gives attackers both patience and speed: they can wait for one newly connected or unpatched device, then concentrate resources when a fresh opening appears.



Continuous pressure

In the first half of 2026, Black Lotus Labs detected nearly **4,000 brute force attacks** and exploit scanners per day on average—reflecting the persistent volume of automated activity probing internet-exposed infrastructure. Credential stuffing and password spraying remain continuous, low-effort methods for gaining initial access, especially when attackers can distribute attempts across large bot populations.

Black Lotus Labs continues to observe automated brute force at scale against Fortinet FortiGate Firewalls, Cisco ASA and Palo Alto GlobalProtect endpoints.

Opportunistic bursts

Dedicated exploit and scanner infrastructure can appear briefly in the days immediately before and after a disclosure. Black Lotus Labs observed this pattern in exploit activity targeting CheckPoint, Citrix NetScaler, and Ivanti Connect Secure services.

In the case of Ivanti Connect Secure, dedicated exploit servers began targeting the services within **24 hours of disclosure** (CVE-2026-10520). Attackers were quick to shift focus onto all Ivanti services in the immediate aftermath of the vulnerability; within **72 hours of disclosure**, more than **100 dedicated servers** began targeting Ivanti services.

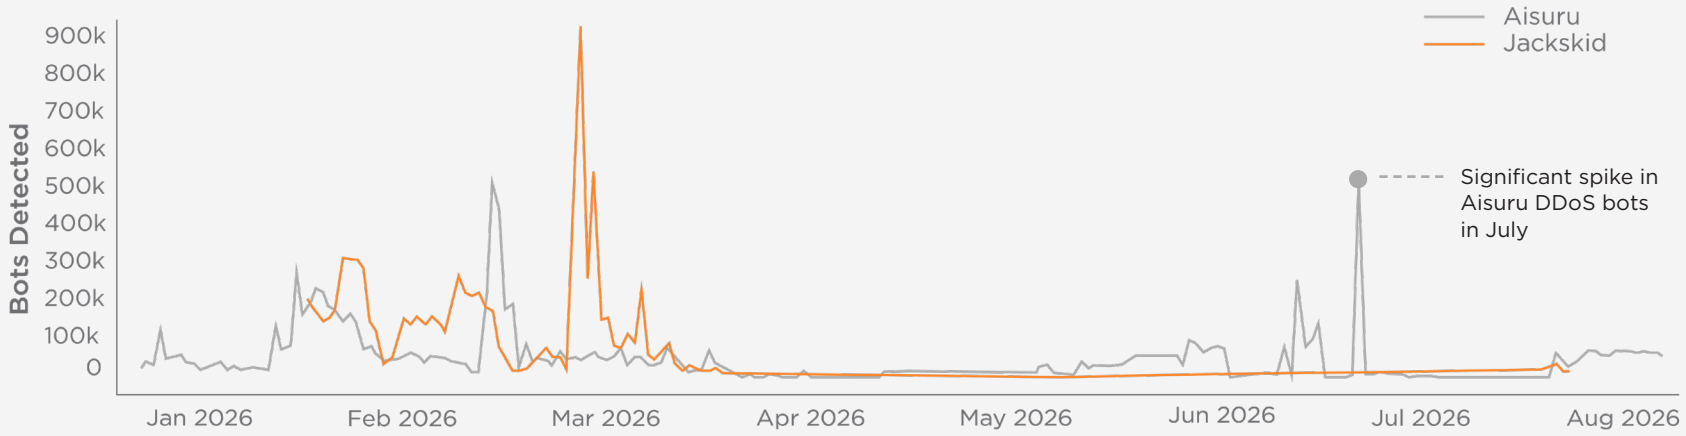
The infrastructure pressure extends to network availability

Attackers' ability to regenerate infrastructure also matters when the objective shifts from access to disruption. DDoS botnets like Aisuru/Kimwolf and Jackskid, which made headlines in early 2026, changed rapidly in size as their operators varied their command-and-control (C2) techniques and rebuilt capacity under pressure.

Despite coordinated action among law enforcement, Black Lotus Labs, and private sector partners to degrade Aisuru/Kimwolf, we observed a significant spike in Aisuru DDoS bots in July. This demonstrates how quickly adversaries can recruit new devices, reconstitute infrastructure, and restore attack capacity. This is the attack-chain regeneration cycle in practice: disruption removes infrastructure, but persistent scanning and bot enrollment continually create the resources needed to rebuild it.

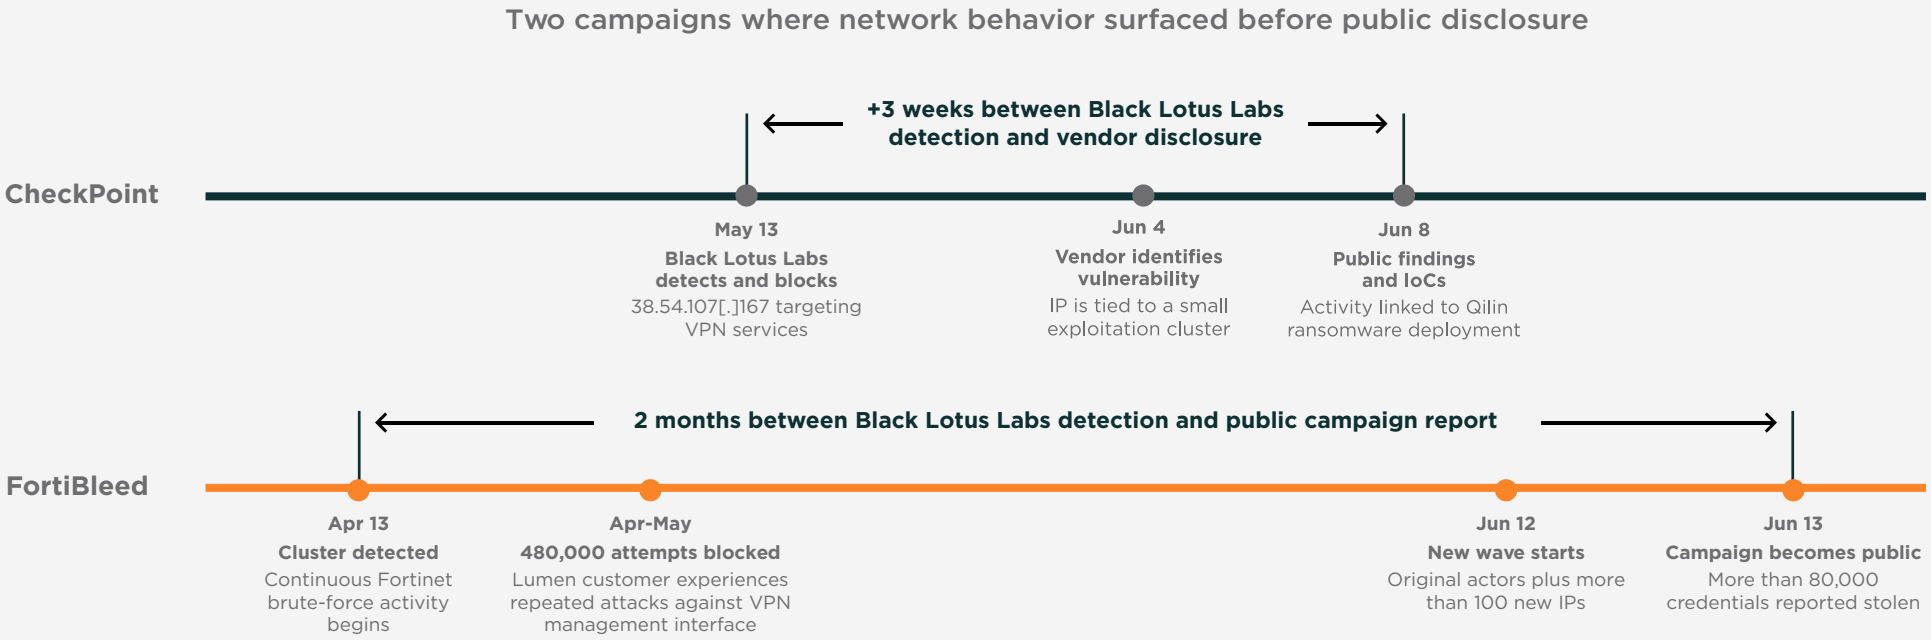
Organizations should assume DDoS capacity can re-form. To help prevent regenerated botnet infrastructure from disrupting operations, teams should continuously validate upstream mitigation, provider escalation paths, and service-availability plans.





The decisive signal may appear before the endpoint alert

Detection and response remain critical for containing activity on managed endpoints. However, conventional visibility misses external proxy pools, unmanaged edge devices, transient scanners, or infrastructure relationships forming outside the enterprise. Network-level intelligence adds a different layer: it can identify coordinated behavior as attackers test services, rotate infrastructure, and converge on the same targets. In one recent campaign, the signal was detected more than **three weeks** before the vulnerability was identified.



CheckPoint: An indicator blocked before CVE disclosure

On May 13, automated Black Lotus Labs detection identified 38.54.107[.]167 communicating with CheckPoint VPN services. The IP was immediately blocked for Lumen Defender customers. On June 4, CheckPoint identified a vulnerability in its own systems and retroactively tied 38.54.107[.]167 to a small exploitation cluster. On June 8, the vendor published indicators, including the same IP detected by Black Lotus Labs, and linked the activity to Qilin ransomware deployment.

What it demonstrates

Campaign-level network behavior can justify protective action before a vendor has completed public attribution. In CheckPoint's case, the block preceded vendor identification by more than three weeks.

FortiBleed: AI-accelerated attacks against exposed infrastructure

On April 13, Black Lotus Labs detected a cluster of IPs repeatedly targeting the same Fortinet endpoints. The campaign continued for more than a month, though the IPs were blocked in Lumen Defender. One customer received more than **480,000 login attempts** against its VPN management interface, all blocked automatically.

Subsequent research found that the FortiBleed actors used the AI-based code editor Cursor and the agentic penetration-testing framework CyberStrike to develop custom tooling and automate intrusion activity. On June 12, the same actor infrastructure began a new wave with more than 100 new IPs. On June 13, a security researcher reported more than 80,000 compromised Fortinet credentials; the activity was later linked to INC/ Lynx ransomware operators.[2]

What it demonstrates

AI can accelerate familiar attack mechanics without fundamentally changing them. NetFlow visibility helps defenders recognize coordinated campaigns as attacker infrastructure pauses, expands, or rotates—and block that activity upstream before it becomes an internal incident.

Lumen internet backbone visibility

99%
public IPv4 visibility

#1
most globally peered network

315B+
NetFlow sessions monitored daily

2.9M
unique threats tracked daily

43K
C2s blocked daily

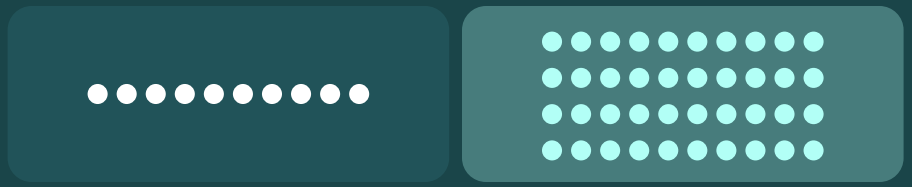
2. "More Than a Leak: What SpyCloud Found Inside the FortiBleed Threat Actor Infrastructure," SpyCloud (June 19, 2026)

Attack infrastructure is becoming regenerative

As we move into the final months of 2026, the most important change is not just that botnets and proxy services are large. It is that recruitment, replacement, and resale are becoming routine operating functions. Compromised devices can be reinfected, reassigned, or offered through multiple services—allowing the same underlying inventory to support several criminal workflows at once.

Regeneration is visible in both recruitment and recovery

Proxy devices contacting tracked malicious APK file hosts increased more than 4x



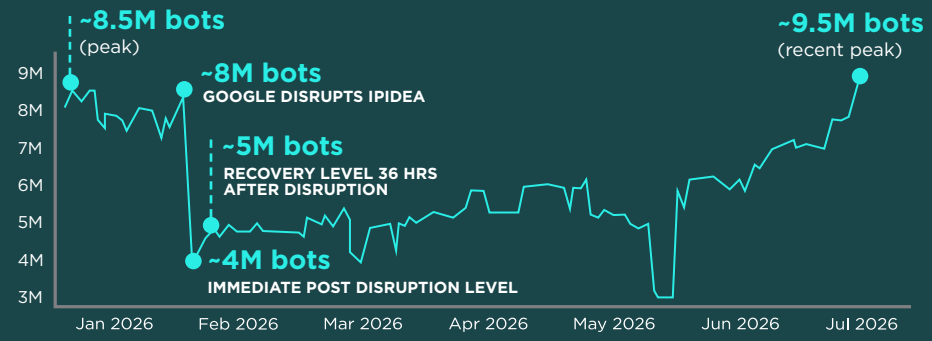
<10k/day
EARLIER IN 2026

→

>40k/day
MID-YEAR 2026

+4x increase in potential proxy devices contacting tracked malicious APK file hosts

IPIDEA disruption and recovery at multimillion-bot scale



+112% Population rebuilt to nearly 112% of pre-disruption levels in less than 6 months

What it means for CISOs

A growing and continually replenished pool of proxy devices gives attackers more infrastructure to rotate through as they scan and target enterprise services—making always-on, network-level visibility and upstream blocking essential to reduce exposure and keep malicious infrastructure from reaching your environment.

Shared inventory expands the adversary supply chain

Among significant botnets Black Lotus Labs tracks, services such as Kookeey, Jaguar, G3Proxy, and PandoraBox have operated at hundreds of thousands or millions of bots. Operators buy and sell IP pools, meaning that a single compromised device can appear in multiple services at the same time. Malicious Android Package Kits (APKs), preloaded proxy software, and vulnerable IoT devices keep adding candidates to the pool.

Earlier in 2026, Black Lotus Labs observed fewer than 10,000 potential net-new proxies per day contacting tracked malicious APK file hosts. By mid-year, the signal exceeded 40,000 per day—a more than **fourfold increase**.

In response, security teams should watch how proxy and botnet infrastructure is probing or targeting the organization's external attack surface, while also looking inward for employee mobile devices, remote-office or retail IoT systems, routers, and other network-connected devices that may be compromised and unknowingly enrolled into these ecosystems.

A device recruited into a proxy network can create a trusted foothold close to, or inside, the enterprise environment. Disruption works, but resilience comes from always-on protection.

Black Lotus Labs has supported action against major botnets like IPIDEA, SocksEscort, ngioweb/NSOCKS, KadNap, and Netnut/Popa, often in collaboration with U.S. and international law enforcement. Some operations ceased observed activity. Others were degraded temporarily. IPIDEA, for example, lost more than **50% of its observed** proxies during a disruption—dropping from a peak of roughly 8M bots in January to 4M. The network recovered quickly, stabilizing at roughly 5 million proxies within approximately 36 hours, in part by reselling proxies. Less than 6 months later, the botnet had rebuilt to a level above its prior tracked population.

The implication here is not that takedowns are futile. Instead, CISOs should assume that malicious infrastructure can reappear, rotate, or be replaced. This means that resilience should be built around always-on preventative protection at the edge and automated blocking that can adapt as proxy and botnet infrastructure changes.

Two supply chains now shape enterprise risk

The adversary supply chain

Vulnerable device → **bot or proxy inventory** → **access and C2** → **scanning, fraud, DDoS, extortion, or ransomware**

Specialization lets operators buy capability instead of building every stage.

The enterprise supply chain

Compromised vendor or service provider → **stolen credentials, poisoned software, or exposed data** → **access to downstream customers** → **data theft, operational disruption, extortion, or ransomware**

One upstream compromise can give attackers access to many organizations and multiply their leverage.

Trend to watch:
Supply chain-based ransomware attacks

Supply-chain compromise is not new. However, the speed and scale with which adversaries can turn trust in a shared tool or platform into access across many organizations—and then hand that access to ransomware or extortion operators—is evolving.

In 2026, a threat actor group known as TeamPCP compromised the development lifecycle for the cybersecurity tool Trivy, creating a poisoned version of the tool which stole credentials from its users. TeamPCP leveraged those credentials to poison additional elements of the cybersecurity/software development ecosystem, including the AI gateway LiteLLM. Credentials were later used for additional exploitation and handed off to multiple ransomware actors—including Vect, who claimed a global provider of financial information, analytics, and market intelligence as a victim. In total, TeamPCP's campaign exfiltrated 300 GB of data, including an estimated 500,000 credential sets.

Separately, the extortion group ShinyHunters has targeted known large enterprise services to maximize reputational harm and notoriety. In March, ShinyHunters alleged it had compromised more than 300 companies' Salesforce portals through a single vulnerability.

In a similar incident in May, ShinyHunters gained access to the Canvas learning management system, a common educational software used by thousands of educational institutions worldwide. In addition to extorting the parent company of the Canvas product, ShinyHunters also attempted to extort the 9,000+ schools whose information was compromised in the breach.

As backups make encryption-only extortion less effective, stolen data and reputational pressure give attackers additional leverage. And because so many organizations depend on the same software, SaaS platforms, and development tools, one upstream compromise can create risk for hundreds or thousands of downstream customers.

CISO implication

Treat software and service dependencies as part of the enterprise attack surface. Maintain visibility into critical third-party software through a software bill of materials (SBOM) and dependency inventories where practical, and extend that visibility to SaaS platforms, continuous integration and continuous delivery/deployment (CI/CD) tooling, API integrations, service accounts, and non-human identities.

H2 2026 defense posture checklist

Organizations must close the gaps between exposure, identity, endpoint, and network layers, then make defensive action as continuous as the pressure being applied.

Priority	Recommended Action	KPIs to the Board
Administrative access	Require MFA, restrict edge management interfaces to controlled paths, and practice least privilege access.	MFA coverage; exposed admin interfaces
Patch exposed systems	Set time-bound SLAs for critical internet-facing vulnerabilities and create an emergency response process for active exploitation.	SLA compliance; time to patch
Retire unsupported edge devices	Identify, isolate, and replace end-of-life (EoL) routers, firewalls, VPNs, and other appliances that cannot be secured.	Percentage of IT-supported assets Risk and cost analysis balanced against alternative protection
Reduce external exposure	Continuously inventory internet-facing assets, remove unnecessary services, and assign an accountable owner to every exposure.	Owned assets; external exposure
Review logs and traffic flows	Centralize authentication, configuration, and management logs. Monitor site-to-cloud and cloud-to-cloud traffic for distributed or low-and-slow activity.	Logging coverage; use cases tested
Correlate control layers	Combine endpoints, identity, cloud, external attack-surface management, continuous exposure management, and network intelligence.	Cross-domain investigations; time to block
Plan for DDoS	Test upstream mitigation, provider escalation, rate limits, and communications before an event. Include botnet-driven and application-layer scenarios.	Exercise date; mitigation readiness
Prepare for extortion and supply chain compromise	Develop response plans for third-party compromise, data theft, supplier outage, and reputational pressure, not just encryption and system restoration.	Tabletop results; decision owners

The bottom line: Cyber risk is increasingly a contest between continuous adversary operations and periodic defensive processes. Security leaders should be able to show that asset discovery, exposure reduction, infrastructure intelligence, and defender response are operating continuously, too.

Accelerating defense to the speed of attack

What we've seen since releasing *The 2026 Lumen Defender Threatscape Report* does not replace the established threat landscape profiled within. Instead, the same operating model is becoming faster, more scalable, and more resilient. Adversaries are applying constant pressure to the edge, concentrating activity around new vulnerabilities, drawing on shared infrastructure markets, and rebuilding operations after disruption. AI may make that sequence cheaper and faster, but the underlying threat vectors are still recognizable.

FortiBleed illustrates what that can look like in practice. Black Lotus Labs detected and began blocking infrastructure targeting Fortinet devices in April, while subsequent research found that the threat actors used an AI-based code editor and agentic penetration-testing framework to develop custom tooling and automate intrusion activity. This shows how threat actors are using AI to develop, adapt, and operate familiar techniques at greater speed and scale.

For CISOs, that changes where the defensive advantage must come from. Trying to outpace machine-speed attackers solely through reactive detection and response creates the wrong contest. Organizations should instead prevent these threats from becoming incidents in the first place. That means strengthening security posture and continuously reducing the external attack surface, while using proactive threat intelligence and network-level visibility to identify malicious infrastructure before it reaches—or successfully exploits—the enterprise.

Going into 2027, the opportunity is to connect these capabilities into a more preventative operating model. Exposure and posture management can identify what attackers are able to reach, while behavior-based network infrastructure intelligence reveals how emerging campaigns are behaving and where malicious infrastructure is forming. Policy-driven automation then turns that intelligence into action—reducing exposure, blocking high-confidence threats upstream, and adapting protections as attacker infrastructure changes.

This, in turn, provides existing detection and response, identity, cloud, and perimeter controls investments with a stronger first line of defense. The organizations that combine strong security fundamentals with continuous posture management, proactive threat intelligence, and automated network-level protection will be better positioned to stop more attacks before they become investigations on a SOC analyst's desk.



Recommendations for Lumen customers

The common thread through this paper is how we at Black Lotus Labs see an opportunity to shift security left in the attack lifecycle. For customers, you're already seeing this proactive protection every day. Still, validating and extending this protection across your network, security stack, and enterprise tooling can deepen the security coverage of the advanced visibility we offer. See below how you can deepen protection against trends from this threat intel in your Lumen security products.

Lumen Defender

- Ask your security team to review Defender for the status of blocked IPs (like those included in the FortiBleed attack) and their mapped threat families to understand the impact of these attacks to your enterprise.
- Your team should use the API to pivot the IPs against publicly available (e.g. AbuseIPDB, VirusTotal) or proprietary threat intel feeds or validate these threats with your MSSP to ensure malicious IPs are blocked across all other edge surfaces in your network (ISP, firewall, etc).
- Vet the impact of Defender's blocked IPs to relief on on-premise appliances to report threat reductions and SOC team time savings to the board.

Lumen DDoS Mitigation

- Review DDoS countermeasures and alert thresholds with your Lumen SOC team to ensure appropriate tuning with business and app-specific requirements.

  **Excludes DDoS Essentials*

Lumen Defender Threat Feed

- Validate with your threat hunting team that profiled IPs from Black Lotus Labs are correlated against your enterprise traffic to spot threat groups that may slip through other points of your network.

Lumen Defender AMDR

- Connect with your AMDR SOC team about exposure management options and ensure unsupported devices are vetted against your exposure management strategy.
- Ask your AMDR SOC team about their response plan for supply chain extortion and align on gaps and policy adjustments.







lumen.com | apac.mail@lumen.com

This content is provided for informational purposes only and may require additional research and substantiation by the end user. In addition, the information is provided “as is” without any warranty or condition of any kind, either express or implied. Use of this information is at the end user’s own risk. Lumen does not warrant that the information will meet the end user’s requirements or that the implementation or usage of this information will result in the desired outcome of the end user. All third-party company and product or service names referenced in this article are for identification purposes only and do not imply endorsement or affiliation with Lumen. This document represents Lumen products and offerings as of the date of issue.

*Lumen Internet On-Demand requires a Lumen Internet port under a minimum term agreement with early termination fees.

© 2026 Lumen. All Rights Reserved.