REPORT

Lumen Quarterly DDoS & Application Threat Report

Q3 2023

LUMEN®

Executive Summary

DDoS

- An overview of AVrecon
- Attack sizes, lengths, and frequencies
- Attack vectors
- Targeted industries

Application Threats

- Blocked application traffic
- Top 5 blocked application suspicions
- Blocked API traffic by industry

The purpose of the Lumen Quarterly DDoS & Application Threat Report is to provide you with an overview of the DDoS and application-layer attacks that are targeting organizations and put them into context. Lumen examined data from the Lumen® DDoS mitigation platform and our API and Application partner, ThreatX, to develop this report.

Don't have time to read the full report? Here's what you need to know at a glance:

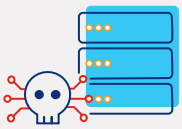

Router malware poses an enormous opportunity for attackers. AVrecon and its 70,000+ compromised devices easily could have been leveraged for DDoS attacks. Learn more.

Bankers beware! The banking industry was highly targeted this quarter with complex, large-scale, multi-vector attacks. Learn more.

Prepare for brute force. Brute force-style application layer attacks are on the rise, with threat actors relying on credential stuffing and programmatic access. Learn more.

LUMEN®

Table of Contents

LUMEN®

Edge vulnerabilities: The implications of AVrecon

What is Black Lotus Labs®?

BLACK LOTUS LABS® by Lumen

Black Lotus Labs® consists of security professionals and data scientists whose mission is to leverage Lumen's global network visibility to help protect our customers and keep the internet clean. Black Lotus Labs uses threat hunting and analysis, machine learning, and automated threat validation to identify and disrupt the work of malicious actors.

Read the latest threat research

Keeping pace with ever-changing cyberthreats is difficult. As quickly as attackers strike with a new tactic, security professionals improve endpoint security to defend against it, then threat actors adapt again in a continuous, high-stakes cycle. One emerging attack surface that saw a lot of action this past quarter was the small-office/home-office (SOHO) router.

What are SOHO routers?

The COVID-19 pandemic redefined the work environment, driving people to hunker down at home and organizations to rapidly develop ways to enable a productive work-from-home experience. This quick shift to remote work involved hastily implemented network infrastructure outside of the relative safety of the corporate datacenter. As people began to rely more and more on their home-office network equipment, threat actors realized a pervasive vulnerability — SOHO routers, which are the weakest points of the network perimeter because they reside at the edge of the network, beyond the corporate firewall.

Why are edge devices vulnerable?

In addition to being physically removed from many of the traditional, on-premises security solutions, consumer edge devices are rarely updated or patched. In addition, many of them are reaching their end of life, meaning that while they may still work, they're no longer supported by manufacturer updates. These vulnerabilities allow threat actors to gain access to SOHO routers with malicious code and then leverage that access to lurk and spy on the target network. This is difficult to detect until an attacker makes a move, so threat actors can hide in plain sight for months or years.

But how does this happen? The processor within SOHO routers — the brain of the device — can be tricked into giving attackers access. Unfortunately, the same processor in your SOHO router could be in your gaming console or thermostat or any other internet of things (IoT) device. This means it's not only SOHO routers that are vulnerable, it's all edge devices, and as the sheer volume of IoT devices increases, the attack surface expands.

AVrecon: A case study

Lumen's threat intelligence team, Black Lotus Labs, has tracked the abuse of networking equipment in various campaigns from nation-state to cybercrime and even hacktivism. This past quarter, Black Lotus Labs identified another multi-year campaign involving compromised routers across the globe. In fact, they discovered one of the largest botnets targeting SOHO routers seen in recent history, dubbed AVrecon.

> “The AVrecon malware operated undetected for more than two years, infecting more than 70,000 machines in that time in more than 20 countries.”

The malware operated undetected for more than two years, infecting more than 70,000 machines in that time in more than 20 countries. This global network of compromised SOHO routers gave cyber criminals the ability to bypass some standard network-based detection tools, especially those based on geolocation, autonomous system-based blocking or IP address-based rate limiting. Once infected, the threat actors leveraged those devices to enable a range of criminal activities. They clicked on various Facebook and Google ads, most likely to commit advertising fraud, and interacted with Microsoft Outlook for suspected password spraying and/or data exfiltration.

Of note, these activities are sneakier than a loud, aggressive DDoS attack, which would automatically expose the infected SOHO routers as part of the botnet and burn the devices for future use. Instead, AVrecon’s operators chose to fly under the radar, focusing on networking equipment that didn’t offer standard endpoint detection and response (EDR) solutions and weren’t likely to be patched against common vulnerabilities.

Black Lotus Labs null-routed the malicious infrastructure, blocking AVrecon across the Lumen network backbone, and added the indicators of compromise (IoCs) from the campaign into the threat intelligence feed that backs Lumen security solutions to block future threat activity. Read more about the team’s detection and mitigation.

Who is at risk?

AVrecon’s operators chose to play the long game, focusing predominantly on stealing bandwidth — without impacting end users — to create a residential proxy service to launder malicious activity. However, AVrecon and its 70,000+ compromised devices could have easily been leveraged for DDoS attacks; once infected, the threat actors can leverage the SOHO devices as they please.

In fact, Black Lotus Labs suspects that hundreds of thousands, perhaps millions, of devices were vulnerable to the AVrecon campaign, and that only the threat actor’s desire for anonymity kept the number of infections to 70,000+. It stands to reason then that if AVrecon’s operators wanted to launch a massive DDoS attack, they would have had a botnet of potentially more than one million devices at their disposal. Even if an attacker used only 50,000 bots at a time to launch DDoS attacks, with a continuous supply for future attacks, the ramifications for victims would be catastrophic. By comparison, in 2016 the infamous Mirai botnet infected hundreds of thousands of IoT devices to produce terabytes of DDoS attacks, taking major websites offline.

Every household contains the kinds of IoT devices leveraged in these attacks, and as the number of IoT devices grows, so does the pool of potential bots for cyberattacks. And with a theoretically limitless supply of leverageable devices, anyone can be a target.

Defending the edge

If any company can be a target, then every company needs to think about its cyber defenses. Consider the following proactive security practices:

- Look for attacks on weak credentials and suspicious login attempts,
- Protect cloud assets from communicating with bots that are attempting to perform password spraying attacks,
- Regularly reboot routers and install security updates and patches,
- Leverage properly configured and updated EDR solutions and regularly update software,
- Ensure all data-in-transit is encrypted and use Secure Access Service Edge (SASE) or similar solutions that use VPN-based access to protect data.

How many DDoS attacks were there?

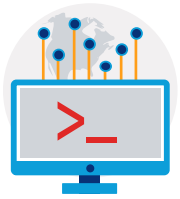

Lumen mitigated

4,217

DDoS attacks in Q3 2023

▼23%

from Q2 2023

51

attacks/day

As with any business, cyberattacks have seasonal ups and downs. Throughout the year, attack trends, styles, techniques and frequency ebb and flow. In Q3 2023, Lumen mitigated 4,217 attacks. This was a 23% quarter-over-quarter decrease and a 24% annual decrease, which could be due to the spike in attacks last year following the invasion of Ukraine. On average we mitigated 51 attacks daily, with September 21, 2023, being our most attacked day in the quarter (298 attacks), followed by July 18 and 4 (153 and 110 attacks respectively). By this time in 2022, Lumen had mitigated more than 16,200 attacks, and at this point in 2023 we've mitigated over 18,300, this is a 13% increase from year over year.

How large are the DDoS attacks?

Largest attack scrubbed

	Dropped bits/s	Dropped pkts/s
Q3 2023	485 Gbps	121 Mpps
Q2 2023	711 Gbps	202 Mpps
QoQ change	▼32%	▼40%
YoY change	▼2%	▼25%

LUMEN®

There are two primary metrics for volumetric DDoS attacks:

Bandwidth attacks: Aim to disrupt service by flooding a circuit or application with traffic. This type of attack is measured in bits per second.

Packet-rate attacks: Consume resources on network elements such as routers and other appliances, as well as servers. These are measured in packets per second with rates typically larger than bandwidth attacks.

In the first half of 2023, Lumen mitigated several large bandwidth attacks (Q1: 817 Gbps and Q2: 711 Gbps), but this quarter we saw a 32% drop in size with the largest bandwidth attack mitigated at 484.96 Gbps. However, we observed the opposite trends with the average bandwidth attack size: we saw an annual increase of 54% from 1.3 Gbps to 2.2 Gbps. The largest attacks (both bandwidth and packet rate) were targeted toward the telecommunications industry, with most attacks hitting around the July 4th holiday weekend.

Bandwidth attacks

- The largest bandwidth attack mitigated in Q3 was 484 Gbps. This was a 32% decrease quarter-over-quarter and a 2% annual decrease.
- The average attack size was 2.2 Gbps, a 12% increase from Q2 and a 54% increase annually.

Packet-rate attacks

- The largest packet-rate attack mitigated in Q3 was 120 Mpps. This was a 40% decrease quarter-over-quarter and a 25% annual decrease.
- The average packet attack size was 320 Kpps, which was a 62% decrease from Q2 and a 26% decrease annually.

How long are DDoS attacks lasting?

Attack duration numbers are affected by the customer's mitigation model. There are two options:

1. **On-Demand mitigation:** Traffic is always monitored, but only scrubbed once a threat has been detected.
2. **Always-On mitigation:** Traffic is constantly scrubbed to further minimize downtime.

The data points in this section only portray trends for On-Demand customers, which account for 69% of attacks mitigated in Q3 2023.

Do I need On-Demand or Always-On mitigation?

	Q2 2023	QoQ change	YoY change
Median attack duration	10m 0s	▲12%	▲5%
Average attack duration	3h 59m 34s	▲91%	▲105%
Longest attack duration	6 days	▼14%	-

The longest attack-period duration we mitigated was six days. It's important to note that this doesn't mean that there was a single attack that lasted six days; rather, it means there was an active campaign which could have contained multiple attacks over time.

We saw a trend of average and median attack-period durations increase over the year. Our average attack duration doubled quarter over quarter and annually, jumping from two hours to an average of nearly four hours in Q3 2023.

Distribution by duration

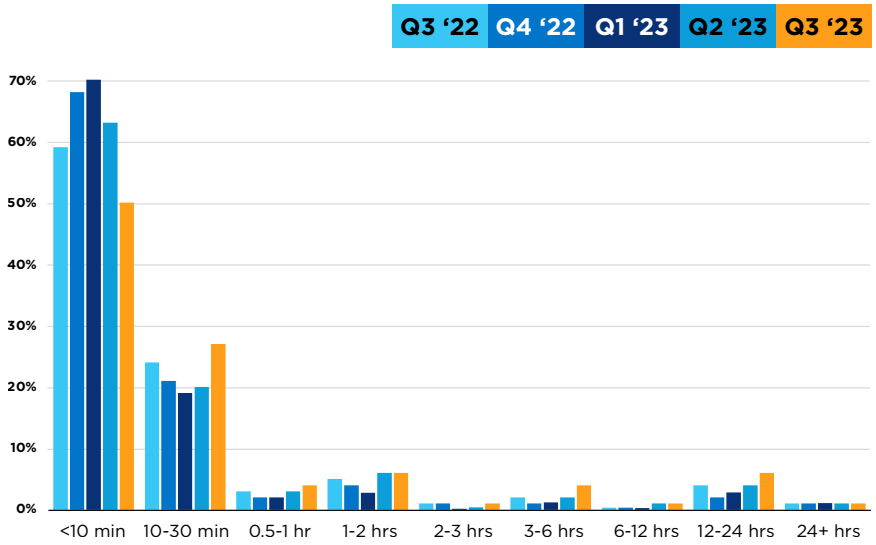

Fifty percent of attacks on Lumen On-Demand DDoS mitigation customers in Q3 were under 10 minutes. This is a 20% decrease from Q2 and a 15% decrease year over year. The second most common attack-period duration was 10-30 minutes, representing 27% of all activity.

> "The average bandwidth attack size saw an annual increase of 54% from 1.3 Gbps to 2.2 Gbps."

Distribution by day

For the most part, attacks were spread evenly throughout the week, with the most popular day for attacks in Q2 being Thursday, which accounted for 20% of activity. This correlated with the most attacked day in Q3, which was September 21, and was targeted against the banking industry.

We expected to see a large volume of attacks over the July 4th holiday weekend. Lumen observed a minor spike on July 6; however, it's not the same level of activity historically mitigated during the holiday in 2022 and 2021. The dip in activity could have been from some recent large-scale takedowns that the government conducted.

Distribution by month

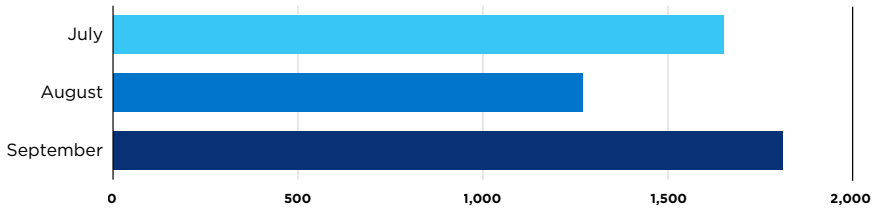

As we zoom out and view Q3 by month we notice that September contained the most attack activity with 1,810 total attacks. August contained the fewest attacks throughout the entire year, with 1,270 attacks. The August activity falls in line with observations from 2022 and shows the seasonality of DDoS attacks.

LUMEN®

What is a multi-vector attack?

Multi-vector attacks are a combination of attack vectors or techniques to compromise a target network or application. Multi-vector attacks are often more sophisticated and may be more difficult to detect than single-vector attacks. They require a higher level of planning and coordination on the part of the attackers, and they often involve multiple stages or phases. As a result, multi-vector attacks can cause significant damage to the target organization, including data theft, financial loss, and reputational damage.

Learn more

What do DDoS attacks look like?

Multi/single-vector attacks

	Q2 2023	Q3 2023	QoQ change
Multi-vector	44%	35%	↓21%
Single-vector	56%	65%	↑16%

In Q3, Lumen observed a shift in attacks for the first time in 2023; multi-vector attacks accounted for 35% of activity and single-vector attacks accounted for 65%. While there was a 21% decline in multi-vector activity, they remained extremely prevalent in one industry: Banking (69% of activity).

Single-vector mitigations

Top single-vector mitigation type breakdown

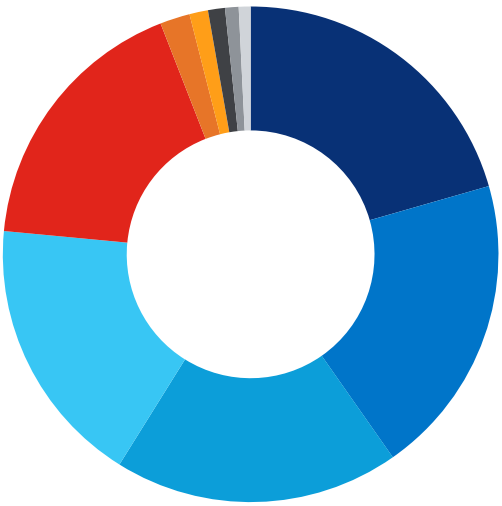

		QoQ
Static filtering	20%	↓15%
Invalid packets	19%	↑3%
DNS	18%	↑25%
UDP	17%	↓8%
TCP SYN	17%	↑4%
Other volumetric	2%	↓9%
per_connection_flood_protection	1%	↑1747%
dns_auth	1%	↑140%
SIP	1%	↓10%
IP fragmentation	1%	↓59%.

For the second quarter in a row, static filtering remained the most common single-vector attack method, accounting for 20% of activity. Static filtering countermeasures are typically done on items such as port and protocol. These statistics also include known bots and abused reflectors as discovered by Black Lotus Labs, which provides initial mitigation against attacks.

Following trends from the first half of 2022, Invalid Packets and DNS mitigations remained popular, 19% and 18% of activity respectively. Last quarter, we observed a short-term spike in UDP amplification, but it decreased 8% quarter over quarter.

Multi-vector mitigations

Top multi-vector mitigation type combinations

		QoQ
DNS, TCP SYN	15%	▲17%
DNS amplification, IP fragmentation, invalid packets, static filtering	9%	N.A..
DNS, TCP SYN, static filtering	8%	▲59%
DNS, static filtering	6%	▲8%
Invalid packets, static filtering	5%	▲118%
UDP, Static Filtering	4%	▲5%
Invalid packets, UDP	4%	▲43%
DNS amplification, IP fragmentation, invalid packets	3%	N.A.
TCP SYN, static filtering	2%	▼9%
Invalid packets, UDP, static filtering	2%	N.A.

DNS amplification combined with TCP SYN was the most used vector combination, accounting for 15% of activity in Q3. Because both methods leverage ports that cannot be turned off or blocked, defending against this combination requires more sophisticated defense tactics than a simple “deny” rule.

Additionally, in Q3 we observed a new four-vector combination: DNS Amplification, IP Fragmentation, Invalid Packets and Static Filtering. The combination represented 9% of activity and was targeted against the banking industry in the campaign that occurred on September 21.

Who is being attacked?

Largest 1,000 attacks by industry

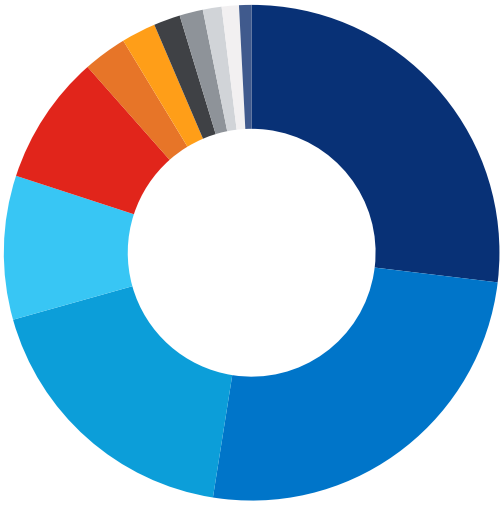

Banking	26%
Government	25%
Telecomm	18%
Software & Technology	9%
Gaming	8%
Finance	3%
Insurance	2%
Utilities	2%
Hosting	2%
Manufacturing	1%
Other	1%
Media & Entertainment	1%
Healthcare	1%
Business Services	1%
Education	0.20%
Retail & Distribution	0.20%

LUMEN®

Of the 1,000 largest attacks Lumen mitigated in Q3, 87% targeted these top five verticals (in order): Banking, Government, Telecommunications, Software and Technology, and Gaming. For the first time, the banking vertical was the most targeted industry. A single banking organization experienced a large DDoS campaign on September 21, and they were hit by 232 attacks in a single day with the median attack duration just under 15 minutes. Despite the number of attacks, the customer experienced no downtime.

We saw a decline in activity targeting the telecommunications industry, which went from representing three-quarters of the largest attack activity to 18%. This is due to a highly targeted telecommunications company shifting its DDoS mitigation strategy.

Banking

- **26%** of the largest 1,000 attacks
- **393** total attacks
- Longest attack period duration: **4 days**
- Largest bandwidth attack: **134 Gbps**
- Largest packet-based attack: **13 Mpps**
- **69%** multi-vector attacks

Government

- **25%** of the largest 1,000 attacks
- **1,683** total attacks
- Longest attack period duration: **6 days**
- Largest bandwidth attack: **8.7 Gbps**
- Largest packet-based attack: **7.3 Mpps**
- **69%** single-vector attacks

Telecommunications

- **18%** of the largest 1,000 attacks
- **542** total attacks
- Longest attack period duration: **6 days**
- Largest bandwidth attack: **485 Gbps**
- Largest packet-based attack: **121 Mpps**
- **61%** single-vector attacks

Software & Technology

- **9%** of the largest 1,000 attacks
- **395** total attacks
- Longest attack period duration: **5 days**
- Largest bandwidth attack: **387 Gbps**
- Largest packet-based attack: **34 Mpps**
- **65%** single-vector attacks

Gaming

- **8%** of the largest 1,000 attacks
- **267** total attacks
- Longest attack period duration: **2 days**
- Largest bandwidth attack: **442 Gbps**
- Largest packet-based attack: **55 Mpps**
- **58%** single-vector attacks

LUMEN®

ThreatX API and Application Protection

As a leader in application and API protection, ThreatX is proud to help protect customers against these threats. Lumen + ThreatX offers comprehensive protection against a wide range of attack types, including programmatic access, command injection, credential stuffing, evasion and more. By leveraging ThreatX's advanced technology, our customers can rest assured their applications are protected against the most common types of attacks.

Learn more by viewing the ThreatX on Lumen Data Sheet

Application protection

Lumen Application Protection uses industry-leading vendors to provide robust application security for our customers. In this report we take a deep dive into the data from one of these partners — ThreatX.

What application traffic is being blocked?

ThreatX requests

	Q2 2023
Total requests	68 Billion
Total blocked requests	417 Million
Blocked bot traffic	205 Million
Blocked due to other reasons	213 Million

Percentage of blocked traffic that was malicious

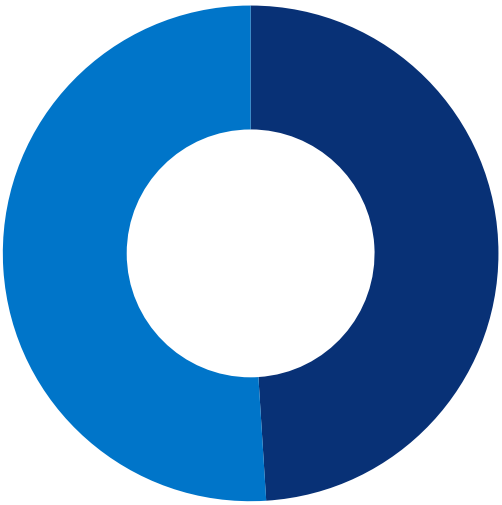

Blocked bot traffic	49%
Blocked due to other reasons	51%

In the third quarter of 2023, ThreatX customers received a total of 68.2 billion requests to their applications. Of these, 417.8 million requests were blocked in real-time, with a fairly even split between blocked bot traffic (204.7 million) and traffic that was blocked for other reasons (213.1 million). These reasons may include specific parameters set forth to access applications or other potential threats attempting to access the applications.

Let's dig into the specific kinds of suspicious activity that ThreatX blocked and which industries were targeted in Q3.

Programmatic access

The ability to manage, control, and interact with a web application or its resources through automated processes, typically by using APIs or scripts, rather than through a user interface.

Credential stuffing

The use of automated tools to try large numbers of username/password combinations, often obtained from previous data breaches, to gain unauthorized access to user accounts.

DirTraversal

DirTraversal, or Directory Traversal, attacks attempt to access files and directories that are stored outside the web root folder.

ErrorRate

Attacks that intentionally induce errors in applications at a rate that overwhelms the system, designed to either exhaust system resources, degrade performance, disrupt service, or expose vulnerabilities.

Evasion attacks

Attacks designed to avoid detection typically with the objective of delivering a malicious payload to a target without being identified.

Top 5 blocked application suspicions

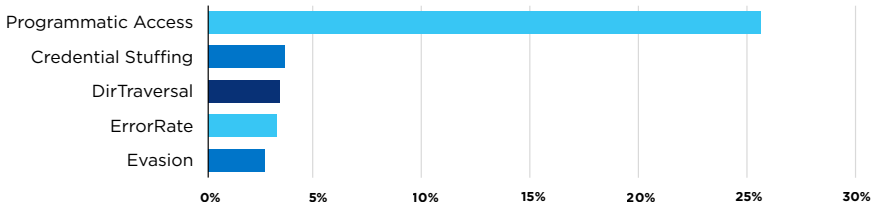

Looking at the traffic ThreatX monitored for customers, some notable and suspicious activity occurred throughout Q3. This includes:

- The highest percentage of blocked traffic came from **programmatic access**, suspicious automated attempts to access a web application, accounting for 25.49% of the blocked traffic. This number is up 89% from Q2.
- Next, **credential stuffing** — which aims to gain unauthorized access to accounts by using known username-password combinations from previous data breaches — accounted for 3.53% of attacks.
- **DirTraversal** and **ErrorRate** accounted for 3.29% and 3.16% of attacks, respectively.
- Rounding out the top five most common attacks are **evasion attacks** at 2.58%.

Interestingly, the top two attack types, programmatic access and credential stuffing, are both types of brute-force attacks that attempt to gain access to resources by trying multiple combinations of credentials. Both attack types also exhibited high variability this quarter, which could indicate targeted campaigns or industries. Brute-force attacks are particularly concerning for sectors with sensitive data, as success could expose valuable personal information, such as bank account details, health information and personal identity details.

Blocked API traffic by industry

Attack ratio by industry

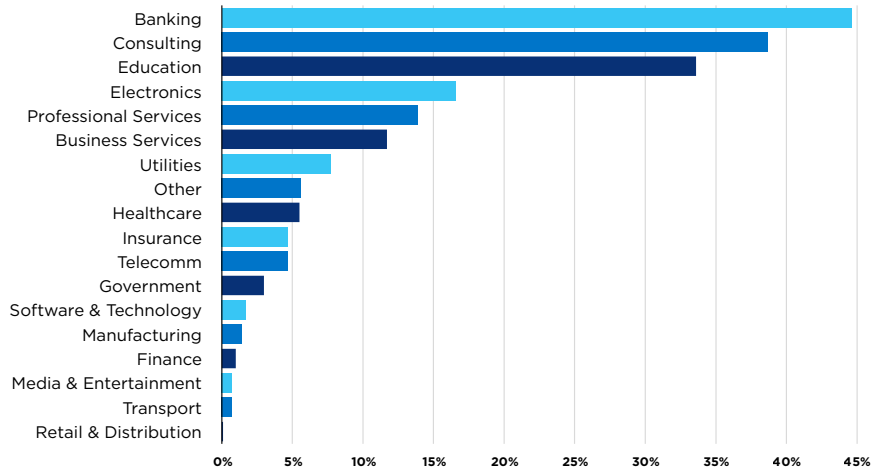

LUMEN®

Looking at the blocked application traffic by industry in Q3 offers a more nuanced understanding of the attack landscape that can point to which industries need tailored security strategies.

High stakes in banking

The banking sector had the highest attack ratio of all industries in Q3, signifying heightened malicious activity and possibly attack complexity. Banking also experienced a significant percentage of Attacks Against Authentication (nearly 25%), which are used to gain unauthorized access to financial data.

Financial institutions are very attractive to attackers, as evidenced by the high attack ratio and combination of brute-force attacks that banking faced in Q3. It's high stakes protecting financial data, but robust web application and API protection solutions can help protect the industry.

Multifaceted attacks in professional services

While this sector doesn't top the list in terms of attack ratios, it does stand out in its variety of attacks and the relatively high number of blocks per customer. It had a relatively high percentage of Bot Attacks (52%) and was the second-most targeted by Attacks Against Authentication (10%). This diversity indicates a complex threat landscape requiring multifaceted defense strategies.

Healthcare and personal data

The Healthcare sector showed a notable percentage of Bot Attacks (55%) and Attacks Against Authentication (6%). This is particularly concerning given the sensitive nature of healthcare data, making it a lucrative target for attackers.

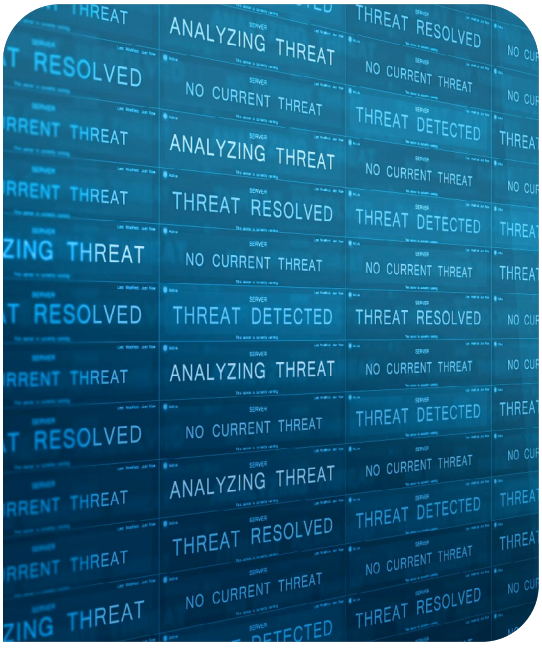

Targeted attacks in insurance

While the Insurance sector may not have the highest attack ratio, it was the third most susceptible to Attacks Against Authentication (7.5%). This means cybercriminals probably took a more targeted approach, aiming to compromise valuable customer data.

Government's unique challenge

The Government vertical had a unique attack type in the top three: Bad Bot, constituting nearly 25% of attacks. This could imply politically motivated cyberattacks or attempts to compromise national security.

By recognizing the common themes derived from these attacks, you'll be able to detect reconnaissance attempts and thwart potential larger-scale API and application attacks down the line. It is important to implement layers of security that harden the application and API and address potential attackers, ensuring that you stay ahead of the game and remain secure.

LUMEN®

Final thoughts from Lumen

Cybercriminals will continue to poke and pry at networks and applications to deny service or exploit data, so you need to be well prepared to stand up to any attacks thrown your way. You never know what the next big attack vector will be, or which industries will be targeted, but one thing is certain: the best defense is to have a solid, in-depth strategy in place.

Recommendations:

- DDoS mitigation is considered basic cybersecurity hygiene. Just like brushing your teeth to avoid cavities, having DDoS mitigation in place can prevent attackers from successfully launching large campaigns against your organization.
- Monitoring your network traffic can help detect if you're under attack, but it can also show if you're being used as a proxy in an attack against someone else. Then, it's a matter of finding, isolating and removing the malware.
- If your company uses applications to interact with customers, employees, or other stakeholders, then you should have holistic protection against network- AND application-layer attacks. This will help ensure your critical business functions stay up and running — even if you are under an active attack. Consider deploying additional application-layer defenses using Web Application Firewalls, API protections and Bot Risk Management solutions, and pair those with application acceleration solutions to make applications more responsive for your customers.
- While the perception is that it's easy to tell if you're under a DDoS attack, tactics are becoming more surgical and discreet. This guide can help you find out if you're under an active DDoS attack.

Hopefully you found this report to be interesting and engaging, and we want to thank you for your time and attention. If you would like to continue learning about trends we have observed, you can read our past quarterly reports.

LUMEN®

Lumen named 2022 Overall Network Security Solution Provider of the Year by *Cybersecurity Breakthrough*.

Read our exciting news

How can Lumen help with DDoS mitigation?

With one of the largest DDoS mitigation deployments in the industry, backed by 170 Tbps of network-based mitigation capacity enacted at more than 500 multi-tiered scrubbing locations, Lumen operates DDoS mitigation at scale. You'll get to choose the mitigation level that is right for your organization with options like On-Demand or Always On mitigation, and advanced features like intelligent scrubbing to help reduce latency and improve performance. You'll also be able to take advantage of a flat monthly service rate. You don't control the length, size or frequency of attacks so why should you be charged for it?

Visit our website to see what DDoS mitigation solution fits you best.

Need immediate protection? Lumen® DDoS Hyper® can be ready in minutes.

Learn more about our advanced DDoS Mitigation Service.

How can Lumen help with application protection?

Lumen Application Protection offers an integrated solution that provides application availability, performance, and security in a DevSecOps-friendly environment for rapid, flexible turn-up of protection against multi-vector and mixed application layer attacks. Lumen partners with a wide variety of Application Protection providers with capabilities spanning web application firewall, bot risk management, and API security to give our customers the optimal selection of features based on their needs.

Visit our website to see which Application Protection solution fits you best.

LUMEN®

Methodology

Data in this report are from the timeframe of July 1, 2023, through September 30, 2023.

Scrubbed attacks are defined as either:

- Incidents flagged by high-level alerts mitigated by the platform, or
- Periods in running mitigations where individual countermeasures are dropping traffic, or
- Events where dropped traffic exceed passed traffic.

Attack vectors or mitigation types are identified either by countermeasures dropping traffic, or misuse types flagged in our flow-based monitoring.

Peaks in the data may be attenuated by how rates are averaged over various time increments.

Data from our Always-On customers are aggregated in increments of minutes, hours or days according to the length of time a mitigation runs. If a mitigation runs long enough that the resolutions time reaches a length of one day, and if there are multiple sequential days of the attack, then it is counted as a single multi-day period of attack.

Data from ThreatX was derived from an analysis of customer traffic.