

Report

# Fastly Threat Insights Report

# Table of Contents

## 03 Executive Summary

## 04 Glossary

## 04 Findings and Insights

04 Network Learning Exchange (NLX)

08 Out-of-Band Domains

10 Bot Traffic

## 15 Recommendations

Fastly's Security Research Team provides the latest attack trends and techniques obtained from traffic signaled on by our Next-Gen WAF. The Next-Gen WAF protects over 90,000 apps and APIs\* and inspects 5.5 trillion requests per month\*\*, covering a wide variety of industries including some of the largest e-commerce, streaming, media and entertainment, financial services, and technology companies around the world.

Our broad reach, across various architectures, including edge and cloud-native environments, allows us to generate insights that are relevant and actionable. These insights can be used to enhance your awareness of web security risks in your industry and organization.

\*As of March 2022

\*\*Trailing 6-month average as of August 1, 2024

# Executive Summary

Our Q2 2024 Fastly Threat Insights Report is focused on highlighting security trends, vulnerabilities, and attack vectors across the web application and API security landscape. With millions of traffic requests across our global customer base, our data provides us with real-time insights into what's materially impacting security teams in the context of larger trends. The goal of this report is to provide security practitioners and decision makers with relevant insights to take proactive measures.

In this report, our security research team builds upon last year's Network Effect Threat Report and other threat publications and provides attack updates and fresh guidance. In addition to the data collected from our Network Learning Exchange (NLX) – the Next-Gen WAF's collective threat intelligence feed – we also examine Out-of-Band (OOB) Domains and traffic signaled by our newest product line, Fastly Bot Management. This expanded report provides a more comprehensive analysis to reflect today's traffic makeup and threat landscape.

This report provides insights from traffic collected during Q2 2024 ('Reporting Period' - April 1 to June 30, 2024). The bot data was taken from the full quarter while the NLX and OOB data was collected April 11 to June 30, 2024.

## Key Takeaways

- **Ephemeral IPs:** 49% of IP addresses added to NLX were listed for just one day, with the average duration being 3.5 days. Attackers use IPs for a short period to avoid detection, highlighting the importance of adaptive security controls that can mitigate varied threats.
- **Mass Scanning:** 91% of attacks originating from NLX sources targeted multiple customers, underscoring how organizations can benefit from collective threat intelligence. Notably, 19% of these sources targeted over 100 different customers.
- **Out-of-Band Domains in CVE exploits:** We observed a consistent rate of exploitation attempts targeting three WordPress plugin CVEs that utilized seven out-of-band domains to inject malicious content, install backdoors, and track infected applications.
- **Human vs Bot:** Approximately 36% of traffic originated from bots, while the remaining 64% came from human users. Notably, one-fourth of bot traffic was from verified wanted bots.

# Definitions

Name	Definition
Network Learning Exchange (NLX)	Fastly's IP reputation feed of potential malicious IPs collected from across our customer base, which can be used to preemptively stop attacks.
Out-of-Band Domains (OOB)	Domains used in attack payloads to establish out-of-band communication channels from vulnerable targets to attacker controlled infrastructure.
Traversal Attacks	An attack that attempts unauthorized access to files, allowing threat actors to reveal sensitive information, modify application data, or enable remote code execution.
Cross-Site Scripting	An attack that injects malicious scripts into web applications that gets executed by the end user's browser.
Verified Bot	A benign bot, distinguished from an unwanted bot, and verified by Fastly, ensuring it can access websites for legitimate purposes (e.g. search engine indexing, uptime status monitoring).

## Findings and Insights

### Network Learning Exchange (NLX)

Fastly's NLX is a collective threat feed included in Next-Gen WAF, used to identify and share potentially threatening IP addresses across customer networks. The efficacy of NLX improves as our network grows, allowing us to observe a larger volume of traffic and analyze attacker behavior more comprehensively.

### Age of NLX IPs

The "age" of an NLX IP address refers to how many days it was listed on the threat feed. By analyzing IP age alongside other data points, we gain a more comprehensive understanding of threats. For instance, if an IP is listed as malicious for an extended period, it might indicate a persistent threat. Alternatively, examining the age of an IP can help in constructing a timeline of an attack, such as initial compromise and an attacker's progression.

As part of Fastly Next-Gen WAF, when the number of attacks from an IP reaches a specified threshold, the IP is added to the NLX feed. Once an IP is added, it remains there for 24 hours unless it meets the criteria again, in which case its expiration time is refreshed.

In Q2 2024, 49% of IP addresses added to NLX were listed for just one day, with the average duration being 3.5 days. The ephemeral use of IPs aligns with common tactics used by attackers to use IP addresses only for a short period to evade long-term detection and reduce traceability.




Figure 1. Percentage of NLX IPs by days observed

There were IPs that persisted on the NLX feed for several weeks and months. Although the percentages were much smaller, they could represent persistent threats and cause disruptions if not addressed.

The coexistence of both transient and persistent threats, underscores the importance of adaptive security controls that can mitigate varied threats.

## Mass Scanning

When you encounter an attack, it's natural to wonder if you're being specifically targeted or if it is a result of a larger pattern. Understanding that you might not be the sole target can help shift the defensive mindset, encouraging organizations to participate in threat intelligence sharing and defense initiatives.

In Q2 2024, 91% of attacks originating from NLX sources targeted multiple customers. Notably, 19% of these sources targeted over 100 different customers. This is a significant increase from [Q2 2023 insights](#), where 69% of NLX sources targeted multiple customers, indicating a rising trend in attacks spreading across broader target bases.




Figure 2. Percentage of NLX IPs by number of customers targeted

Some attackers are not focused on a specific target but are looking for any vulnerable system. Scanning multiple targets simultaneously is more time-efficient, allowing attackers to cast a wider net, increasing the chances of finding systems they can exploit with minimal effort.

## Attack breakdown by industry

Examining NLX data from a broader perspective revealed that attackers targeted the High Tech industry the most, accounting for 37% of attacks. This is a decrease from Q2 2023, where the High Tech industry faced 46% of attack traffic.




Figure 3. Industries ranked by the percentage of targeted attacks

High-tech companies can be particularly attractive targets for attackers due to their broader impact. If compromised, it could enable attackers to extend their attack to all downstream users and organizations that rely on the affected company’s products and services. For example, in June 2024, a breach at Snowflake, a data management company, led to the [exposure of credentials](#) belonging to a number of its customers.

Pivoting on this data and examining the percentage of attacks with and without NLX across industries, the Education industry experienced 22% more attacks originating from NLX sources compared to those without. In other industries, while attacks originating from NLX sources were fewer, they represented nearly 40% of attack traffic per industry on average, highlighting the importance of using a collective approach to IP reputation intelligence.




Figure 4. Percentage of NLX versus non-NLX tagged IPs by industry

### Trending attack techniques

Traversal techniques are the favored attack among NLX sources, constituting 35% of attacks tagged with NLX in Q2 2024. It remains the top attack technique, with a slight increase from 32% in Q2 2023.

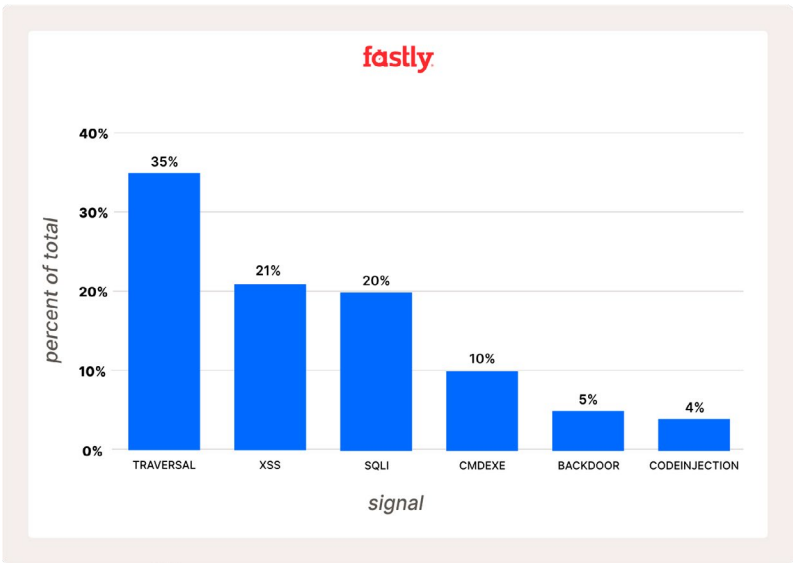


Figure 5. Percentage of total attacks by attack signal

The emphasis on traversal attacks highlights a concern organizations need to be aware of. This is reinforced by the [joint alert](#) by the Cybersecurity and Infrastructure Security Agency (CISA) and the FBI, stating that recent campaigns have exploited traversal vulnerabilities to compromise critical infrastructure in healthcare and other industries.

## Out-of-Band Domains

As part of Fastly Next-Gen WAF, we analyze out-of-band domain usage across all attack payloads. Using out-of-band domains is a common technique used to discover and exploit vulnerabilities, bypass security controls, and exfiltrate data. Monitoring and tracking their usage can help identify compromised systems, detect ongoing attacks, and uncover attacker methodology.

### Active exploitation of unauthenticated stored XSS vulnerabilities in WordPress Plugins

In Q2 2024, [we saw](#) a dramatic increase in usage of several domains related to active exploitation of three WordPress Plugin CVEs (CVE-2024-2194, CVE-2023-6961, and CVE-2023-40000). The exploitation attempts we observed sent various XSS payloads that utilized out-of-band domains to inject malicious content, install backdoors, and track infected applications. Exploitation attempts originated from thousands of IP addresses, with a geographic concentration in the Netherlands, but utilized only seven different domains throughout the campaign. Once the exploitation activity began, it remained at a consistent rate throughout the quarter, as seen in the provided graph.




Figure 6. Attack volume of Wordpress Plugin CVEs over time



## Out-of-band domain tooling

Typically, the out-of-band domains we observe are part of commercial security scanners and common attack tools. At 53%, the largest portion of out-of-band domains are self-hosted callback servers using open source projects, such as Project Discovery's interactsh and Matthew Bryant's XSS Hunter Express. Invicti Security's tool suite includes both what was formerly Netsparker and Acunetix at 27%. Project Discovery rounds out the top three with almost 10%, noting that this usage is of the domains they control and does not include self-hosted instances of their open source software.

If we exclude the domains used in the Wordpress CVE campaign, this table shows which scanners and tools were used as a percentage of total out-of-band domains signaled.

Out-of-Band Tool	Percent
Other/Self-Hosted	53.06%
Invicti Security	27.88%
Project Discovery	9.97%
Qualys	5.44%
Nessus	3.33%

Figure 7. Distribution of tooling types by percentage of OOB domains

## Out-of-band domain attack types

Looking at the top attack signals that use out-of-band domains, command injection (CMDEXE) has the most at 45%, followed by cross-site scripting (XSS) with 26%.

Attack Signal	Percent
CMDEXE	45.40%
XSS	26.40%
SQLI	13.69%
CODEINJECTION	13.22%
TRAVERSAL	1.22%

Figure 8. Percentage of attack signals using OOB domains

Although path traversal is the most utilized attack type in NLX, it constitutes only a small percentage of attacks involving OOB domains. Path traversal vulnerabilities primarily involve accessing or reading files from the local file system of the vulnerable server and often don't involve outbound network communication, which is essential for OOB techniques.

## Vulnerability Graveyard: Log4Shell

On December 6, 2021, CVE-2021-44228, colloquially known as Log4Shell, took the world by storm. We initially addressed the vulnerability in Fastly Next-Gen WAF by releasing a CVE virtual patch. Acknowledging that organizations would be dealing with this vulnerability on a longer term basis, we later released an improved detection as the LOG4J-JNDI system signal.

Now, more than two and half years later, we made the decision to exclude LOG4J-JNDI signal from our out-of-band domain analysis. We still see scanning activity that tries to exploit this vulnerability, and Fastly Next-Gen WAF will continue to signal on attacks. But due to the age of the vulnerability we determined it is no longer relevant to include in trends.

## Bot Traffic

A significant portion of the internet traffic, as evident from our research, can be attributed to requests generated by automation tools, commonly referred to as bots. Fastly uses a variety of techniques (e.g., client and network analysis, advanced challenges, behavioral analysis) to distinguish a human user from a bot.

While a large portion of this automated traffic is malicious or undesirable in nature, such as account takeover attacks, ad fraud, carding, and others, there are several benign or even beneficial use cases for bots. These use cases range from search engine crawlers indexing a website to uptime monitoring tools. Website owners would want these crawlers and monitoring agents to be able to access the website, while blocking the unwanted ones.

Unfortunately, unwanted bots often impersonate well-known benign bots to evade simple detection and carry out their intended activity. Fastly has curated a list of such well known wanted bots, along with the means to be able to distinguish them from an imposter bot. The requests from these verifiably wanted bots are tagged with the **VERIFIED-BOT** signal in Fastly Bot Management, and they are further classified into various categories based on the main purpose of the bot.

# Bot traffic classification




Figure 9. Distribution of traffic by type

In our analysis, we identified approximately 36% of the traffic to Fastly Bot Management customers originated from bots, while the remaining 64% came from human users. Notably, 75% of this bot traffic was classified as unwanted, whereas the remaining 25% consisted of verified wanted bots.




Figure 10. Distribution of traffic volume by type over time




Figure 11. Percentage of wanted versus unwanted bot traffic by industry

Examining the breakdown of wanted and unwanted bot traffic by industry, Financial Services has a large proportion of unwanted bot traffic. Financial services applications often handle sensitive data such as financial transactions and personal identification information, which make them lucrative targets for credential stuffing, account aggregation, web scraping and other automated attacks by bots.

## Wanted Bots category breakdown

As previously mentioned, wanted bots are categorized based on their intended use case, as described in the following table. A significant portion of wanted bot traffic – nearly 70% – was attributable to “Search Engine Crawlers.” This isn’t surprising given the periodic crawls, extensive scope (crawling almost all unauthenticated pages of a website), and the increasing number of search engines and crawlers on the web.




Figure 12. Distribution of wanted bot traffic volumes by category

# Wanted Bot Categories

Category	Description
Search Engine Crawler	Tools which access your site to show a preview of the page, in other online services, and social media platforms.
Research	Tools which access your site to monitor performance, uptime, proving domain control, etc.
Page Preview	Tools which access your site to show a preview of the page, in other online services, and social media platforms.
Monitoring & Site Tools	Tools which access your site to monitor performance, uptime, proving domain control, etc.
Search Engine Optimization	Tools that analyze page content for SEO purposes.
Content Fetcher	Tools which extract content from websites to be used elsewhere.
Security Tools	Security analysis tools to inspect your site for vulnerabilities, misconfigurations and other security features.
Accessibility	Tools which make content accessible, such as screen readers, etc.
Platform Integrations	Integration with other platforms by accessing the website's API, notably WebHooks.
Online Marketing	Crawlers from online marketing platforms to aid in Ad placement.

## Wanted Bot breakdown by industry

Unlike the overall trend, where the majority of wanted bot traffic was attributed to Search Engine Crawlers, this was not the case for Financial Services organizations. When analyzing the traffic share of wanted bots by industry, Financial Services primarily consisted of Page Preview and Monitoring & Site Tools.

Due to the sensitivity of financial applications, it's plausible that the majority of their content is behind authenticated web pages, making them inaccessible to crawlers. Alternatively, these companies may deliberately configure their sites to instruct search engine crawlers not to scan their sites or to limit the URLs that the crawlers can access.




Figure 13. Percentage distribution of wanted bots by industry

While bot management solutions help keep unwanted bots at bay, it is equally important for website owners to monitor and manage wanted bots visiting their sites. For instance, a “Page Preview” bot can drive traffic to your site, whereas a “Content Fetcher” bot that doesn’t credit the source can have negative effects. Similarly, search engine crawlers increase the visibility of your content, but some may be too aggressive, costing you dearly in compute and bandwidth. As a result, website owners should have tooling that can selectively allow or deny specific wanted bots based on their behavior and alignment with the website’s business objectives.

# Recommendations and Actionable Guidance

The threat intelligence insights presented in this report can help you fine-tune your security strategies. We recommend that Fastly Next-Gen WAF customers use NLX through the **SIGSCI-IP** signal to improve their detections. This signal can be used in conjunction with attack and anomaly signals, along with other custom conditions in rules. You can learn more about using signals in the Next-Gen WAF [documentation](#).

As a best practice, we recommend you monitor and track the usage of out-of-band domains to help identify compromised systems, detect ongoing attacks, and uncover attacker methodology. Fastly Next-Gen WAF customers can take action on threats associated with OOB domains through the **OOB-DOMAIN** signal.

Bot management is an extremely important capability with the ability to monitor and manage both wanted and unwanted bots. This is crucial to preventing automated attacks and other unwanted access to the website, while maintaining access for benign and essential bots. Next-Gen WAF offers the flexibility for customers to control access for individual wanted bots using the **VERIFIED-BOT** signal.

It's important to know that today's web security threats come in various forms, such as web application attacks and malicious bots. We strongly suggest that you adopt advanced mitigation strategies for these threats to comprehensively reduce your security risks.