

Report

Q1 2025 Fastly Threat Insights Report

fastly®

Table of Contents

03 Glossary

04 Findings and Insights

04	NGWAF Attack Trends
08	Network Learning Exchange (NLX)
10	Bot Traffic Analysis
16	Distributed Denial of Service (DDoS) Trends

The Q1 2025 Fastly Threat Insights Report highlights security trends, attack vectors, and threat activity across the application security landscape. Drawing from trillions of requests across our global customer base, this report offers a real-time view into what's materially impacting security teams in the context of larger trends.

This quarter's insights are derived from traffic analyzed across Fastly's Next-Gen WAF (NGWAF), Bot Management, and DDoS Protection products. These solutions collectively protect over 130,000 apps and APIs* and inspect more than 6.5 trillion requests per month**. Fastly's broad visibility spanning edge and cloud-native architectures, combined with our presence across a wide range of industries, including leading e-commerce, streaming, media and entertainment, financial services, and technology organizations, provides us with a unique and comprehensive view of the global web application threat landscape.

**As of April 2025*

***Trailing 6-month average as of April 2025*

Key Takeaways

1. Fastly's Bot Management data revealed that over one-third (37%) of all observed traffic came from bots, while 63% originated from human users.
2. Attacks on the commerce industry doubled, rising to 31% of all observed attacks in Q1 2025 from 15% of all observed attacks in Q1 2024.
3. In Q1 2025 Cross-Site Scripting (XSS) is the most prevalent web attack, increasing from 21% in Q1 2023 to 35% in Q1 2024, and reaching 40% in Q1 2025.
4. 28% of all observed attacks originated from IPs listed on Fastly's NLX, a shared, real-time threat feed of confirmed malicious IPs.
5. Attempted logins using compromised passwords averaged over 1.3 million per day, driven in part by the use of proxy services to automate activities.

Definitions

Name	Definition
Attack Signals	Attack signals are tags applied to malicious requests that contain attack payloads, as defined in our Next-Gen WAF (NGWAF) documentation .
Account Takeover (ATO)	A type of attack to gain unauthorized access to a user's online account through various means, but typically using stolen login credentials.
Bot Traffic	Any non-human internet traffic can be beneficial (e.g., search engine crawlers) or malicious (e.g., carding).
Distributed Denial of Service (DDoS)	An attack that aims to make a website or service unavailable to legitimate users by overwhelming it with traffic.
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.
Points of Presence (PoP)	Strategically located data centers around the world, where Fastly deploys its edge servers for caching content and processing requests closer to end users, reducing latency, improving performance, and making real-time security decisions.
Web Application Attacks	Techniques and methods attackers use to exploit vulnerabilities in web applications and APIs

Findings and Insights

NGWAF Attack Trends

Web applications are constantly targeted by a wide range of attack techniques, from simple probing to sophisticated attempts. NGWAF helps defend against these threats by applying signals to malicious requests that contain attack payloads.

These signals allow security teams to understand not just whether a request is malicious, but what kind of attack is being attempted. By aggregating this signal data, we are able to describe broad trends in attacker behavior.

In this report, we focused on the leading attacks surfaced through NGWAF attack signals. These attacks represent the most commonly observed threats during the analysis period. The table below provides a definition of each attack type.

Attack Types


Cross Site Scripting (XSS)	A vulnerability that allows attackers to inject malicious scripts into trusted websites.
Traversal	A vulnerability that allows attackers to access files and directories outside the intended web root directory.
SQL Injection (SQLi)	A vulnerability that allows an attacker to inject SQL code into web applications to view or modify a database.
Command Execution (CMDEXE)	A vulnerability that allows an attacker to execute arbitrary commands on the host operating system of a vulnerable application.
Backdoor	A technique that allows attackers to bypass security controls to gain unauthorized access to a system.

Attack Trends

In Q1 2025, XSS was the leading attack type, accounting for 40% of all attacks observed across Fastly customers. This continues the upward trend observed in Q1 2024, where XSS comprised 35% of all attacks. In contrast, SQLi, which was the leading attack type in Q1 2023 at 30%, has steadily declined, dropping to 24% in Q1 2024 and 18% in Q1 2025, highlighting a noticeable shift in the threat landscape (Image 1).

The decline in SQL injection attacks could reflect the effectiveness of modern development frameworks and security best practices. Features such as prepared statements, object-relational mapping (ORM) libraries, and cloud-managed databases reduce the risk of SQLi when used correctly, making these attacks harder to exploit and prompting attackers to focus their efforts elsewhere.

Meanwhile, XSS attacks are often low-cost but high-impact. They are ideal for stealing session tokens, injecting malicious code, or harvesting other sensitive information retained by the browser. Modern sites integrate dozens of third-party libraries, analytic scripts, and marketing trackers, often introducing blind spots that can serve as an ideal vector for XSS.



XSS in Focus

To better understand the threat activity driving XSS attacks, we analyzed exploitation attempts. We continue to observe a high volume of attacks targeting unauthenticated stored XSS vulnerabilities in WordPress Plugins – similar to the activity we [reported in 2024](#).

The attacks inject a script tag referencing an obfuscated JavaScript file hosted on an external domain. The scripts used in these exploitation attempts are identical, suggesting a coordinated campaign, focusing on the following malicious actions:

1. Creating a new administrator account
2. Injecting backdoors
3. Setting up tracking scripts, presumably to monitor infected sites

Indicators of Compromise (IoCs) currently observed in association with these exploitation attempts:

Domains:

- gll.instantcontentflow[.]com
- idc.cloudiync[.]com
- cloud.cdndynamic[.]com
- metrics.gocloudmaps[.]com
- cloud.swiftstreamhub[.]com

IP Addresses:


- 93.174.93[.]2
- 89.248.169[.]52

What makes this campaign stand out is its lack of sophistication in covering its tracks. Despite the volume, the attacks are executed from just a couple of IP addresses, where a more distributed approach would typically be used to obfuscate the origin. Even more brazen, two of the domains identified in our original report are still actively in use, and the user agent string associated with the IP address sending the most traffic is self-advertising curl. This threat actor appears unconcerned about hiding their activity.


To detect ongoing attack campaigns, uncover attacker methodology, and mitigate similar activity, we recommend monitoring for attacker-controlled domains embedded in attack payloads. For Fastly customers, the out-of-band (OOB-DOMAIN) signal can be leveraged to surface these domains and take proactive action against emerging threats.

Attacks by industry

In Q1 2025, High Technology organizations were targeted the most, accounting for 35% of all observed attacks. While this represents a significant share, it is a notable decline from 54% in Q1 2024. Meanwhile, the Commerce industry saw an increase, rising to 31% in Q1 2025 from 15% in Q1 2024. Financial Services saw a decline from 17% to 12%, Media & Entertainment from 9% to 14%, and Education from 3% to 4%.



The increase in Commerce could reflect adversary interest in more immediate financial rewards, such as credit card data, PII, and transactional manipulation. A particularly telling data point is that 54% of attacks represented in the Commerce industry are attributed to XSS (Image 3). This aligns with XSS's broader prevalence across all sectors.



The decline in High Tech does not necessarily imply reduced threat. High-tech companies are particularly attractive targets due to their potential for widespread downstream impact. A successful compromise can provide attackers with access not only to the company itself, but also to the many customers, partners, and services that depend on its infrastructure.

A recent example is the [GitHub supply chain attack](#) involving the compromise of the popular GitHub Action `tj-actions/changed-files`, used by over 23,000 repositories. A threat actor added malicious code that prints out secrets in project build logs. This incident demonstrates the disruption and effects attackers seek in targeting high-tech companies.

Adding to this concern is the unusually high percentage of backdoor attempts (15%), compared to minimal or nearly absent levels in other industries. Backdoor attacks are typically used to establish persistent access into systems and are precursors to long-term, multi-stage intrusions. The disproportionate use of backdoor attempts aligns with attacker interest, gaining sustained footholds that serve as an avenue for widespread downstream impact.

Network Learning Exchange (NLX)

NLX is a real-time threat intelligence feed included in NGWAF that shares confirmed malicious IP addresses across customer environments. When an IP exceeds attack thresholds, it is flagged, added to NLX, and automatically shared with a default 24-hour expiration. If ongoing malicious activity is observed from the same IP on the same site or others, the expiration is extended based on the most recent activity. Leveraging NLX enables customers to block threats before they reach their networks and shift from reactive defense into proactive protection.

In Q1 2025, 28% of attacks originated from IP addresses listed on NLX, and 48% of those IPs targeted multiple customers. This data suggests attackers operate opportunistically, casting a wider net rather than focusing on a specific target to increase their chances of finding systems they can then exploit.


28%

Percentage of Attacks
Signaled with NLX

2.9


Avg. Number of Days
IPs Remain on NLX

Notably, 47% of IP addresses listed on NLX were active for just one day, with an average lifespan of 2.9 days, reflecting a common tactic of using short-lived IPs to avoid long-term detection and reduce traceability (Image 5). By analyzing the “age” of these IPs alongside other indicators, we can gain deeper insights into attacker behavior. For example, if an IP is listed as malicious for an extended period, it may signal an ongoing campaign. Moreover, examining the age of an IP can help in constructing a timeline of an attack, such as initial compromise and an attacker’s progression.



NLX by industry

In Q1 2025, industry-level traffic analysis revealed that Education organizations saw the highest proportion of attacks associated with IPs listed on NLX (61%), than attacks from IPs not listed on NLX (Image 6). This suggests that adversaries targeting this sector are often reusing infrastructure that has already been identified and tracked across Fastly's platform.



Other industries also experienced high levels of NLX traffic, including:

- Healthcare (46%)
- Media & Entertainment (41%)
- High Technology (31%)

These numbers reflect the broad applicability of NLX as a cross-industry threat intelligence asset. Rather than relying solely on reactive detection, organizations in these industries benefit from Fastly's ability to proactively flag malicious IPs. This preemptive visibility allows defenders to focus on high-confidence signals and reduce noise from repeat offenders.


Bot Traffic Analysis

A significant portion of the internet traffic is generated by automation tools, or bots. Fastly uses techniques such as network analysis, behavioral analysis, and advanced challenges to differentiate human users from bots.


While a large portion of bot traffic is malicious, ranging from account takeover attempts, ad fraud, carding, and others, there are also legitimate use cases, such as search engine crawlers or uptime monitoring tools where Website owners want to allow beneficial bots while blocking the unwanted ones.

More recently, a new class of bots has emerged in the form of AI bots, which crawl websites either to train large language models (LLMs) or to enrich model responses with grounding at inference time. Whether these bots are seen as a benefit or a risk depends on the site owners priorities. We'll be sharing more detailed insights into these AI bots and their behavior in future reports.

In our analysis, Fastly Bot Management data revealed that approximately 37% of all observed traffic originated from bots. Of that, 89% are unwanted bots, meaning that approximately 33% of total traffic provides no business value and may pose risks like fraud, scraping, and infrastructure strain. This means that as an organization, if you're not actively managing bot traffic, 1 in every 3 dollars you're spending on infrastructure, bandwidth, or performance might be being wasted on serving malicious or non-productive traffic.



Examining the breakdown of wanted and unwanted bot traffic by industry, Commerce websites attract the largest proportion of unwanted bot traffic at 39% (Image 8). This trend aligns with broader attack patterns highlighted earlier in this report. Commerce websites are lucrative targets for cybercriminals to steal sensitive data (like credit card info), exploit vulnerabilities, take over accounts, scrape prices or disrupt operations - often for profit, fraud or even competitive advantage.




Wanted Bots

Fastly maintains a curated list of such well known wanted bots, along with the means to be able to distinguish them from an imposter bot and verifies them with a VERIFIED-BOT signal. These bots are further classified into various categories based on the main purpose of the bot, as described in the following table.

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.

A significant portion of wanted bot traffic (66%) was attributable to Search Engine Crawlers (Image 9). This isn't surprising given the periodic crawls, extensive scope (crawling almost all unauthenticated pages of a website) of search engines and crawlers on the web.



Wanted Bots by Industry

Unlike the broader trend, where Search Engine Crawlers represented the majority of wanted bot traffic, Financial Services organizations primarily consisted of Content Fetcher and Page Preview (Image 10).



Given the sensitivity of financial applications, it's likely that the majority of their content is behind authenticated web pages, making them inaccessible to crawlers. Additionally, these companies may intentionally configure their sites to block or restrict access to limit which URLs can be indexed.


While blocking unwanted bots is essential, managing the behavior of legitimate bots is equally important. Some, like Page Preview bots, can drive engagement, while others, such as Content Fetechers, may undermine content value by not crediting the source. Even search engineer crawlers can strain resources if not properly managed.

To protect performance and align with strategic goals, we recommend that website owners have tooling that can selectively control bot access based on their value and impact.

Account Takeover (ATO) activity

In recent years, the frequent occurrence of data breaches and credential dumps has become an unfortunate reality. Due to credential reuse, cybercriminals leverage compromised credentials to carry out account takeover (ATO) attacks, often using specialized automation tooling or bots to do so at scale. Following the general availability of our compromised password signal in early March, we have gained visibility into credential-based threats impacting customers.

In March 2025, for customers using the compromised password signal, attempted logins with compromised passwords averaged over 1.3 million per day. Notably, there were significant spikes in activity from March 7-10, March 18-20 and again on March 31 (Image 14).



The larger spikes corresponded with coordinated credential stuffing and brute-force campaigns. The high volume, frequency, and diversity of IP addresses observed during these events revealed the use of proxy services to automate and distribute attack traffic.

For example, one coordinated campaign we observed unfolded over a 16-hour period, during which all requests carried a French (fr) language header and a user-agent identifying as Chrome 96 on Windows 10. The campaign generated approximately 1.5 million compromised password attempts from more than 275,000 unique IP addresses, averaging 1.5 to 1.7 attempts per IP per hour. The top source countries were Brazil (15%), United States (11%), Bangladesh (6%), Taiwan (5%), and Vietnam (4%).

Whether you're a user, developer, or security professional, here are steps you can take to prepare.

For Individual Users:

- Use a password manager to generate strong, unique passwords.
- Enable MFA (multi-factor authentication) everywhere possible.
- Check your email on Have I Been Pwned to see if it's been involved in a data breach.
- Be cautious of reusing passwords - reused credentials are a key target in stuffing attacks.

For Site and App Owners:

- Implement rate limiting and IP reputation scoring.
- Monitor for unusual login patterns, such as high failure rates from diverse IPs.
- Integrate with threat intelligence feeds to block known bad IPs.
- Use bot mitigation tools.
- For Fastly customers, integrate compromised password detection into your authentication workflows using the COMPROMISED-PASSWORD signal through NGWAF templated rules.


As the compromised password signal reaches broader adoption, early observations are already revealing valuable insights into credential based threats. We look forward to sharing deeper trends and developments in future reports.

Distributed Denial of Service (DDoS) Trends

Application DDoS attacks are a significant threat to any internet-facing application or API, capable of disrupting service performance and availability for end users and potentially leading to revenue loss for organizations. These attacks specifically target Layer 7 services, such as web applications, with the intent of exhausting server resources with a relatively low volume of cleverly crafted HTTP requests.


Unlike network DDoS attacks, which aim to overwhelm network infrastructure with massive amounts of traffic often measured in terabits per second (Tbps), application DDoS attacks exploit weaknesses in application code or protocol implementation to disrupt services with comparatively less traffic. This can make them harder to detect but no less dangerous.

Since the public release of Fastly DDoS Protection in October 2024, we are observing a slow but steady increase in DDoS attack volumes over the past few months (Image 16).




We also noticed a relatively sizable campaign over a period of 4 days in the middle of March 2025.

DDoS attack traffic was detected at multiple PoPs worldwide, with the highest volumes observed at our PoPs in the US (26%), followed by Germany (18%), Singapore (14%), France (7%), and the UK (6%) (Image 17).



While the attacks predominantly targeted the media and entertainment industry (Image 18), we observed a notable concentration (48%) on small and medium sized businesses (SMBs)—defined as companies with less than \$100 million in annual revenue, in contrast to Commercial (\$100 million-\$1 billion) and Enterprise (over \$1 billion) segments—during this period (Image 19).



Smaller organizations can often operate with limited cybersecurity resources, making them appealing to attackers. Their reduced ability to detect, mitigate, and recover from threats increases their risk exposure. These insights highlight that organizations of all sizes are subjected to application-layer DDoS attacks. It’s critical that even smaller businesses adopt security measures to strengthen their defense and ensure operational resilience.

While large-scale DDoS attacks grab headlines, smaller events occur regularly and can still disrupt availability. We recommend that you ensure your infrastructure and security tools can scale effectively to handle unpredictable volumes of traffic.

About Fastly

Fastly is the application security leader and edge cloud platform behind many top digital experiences. Fastly helps organizations deliver and secure online experiences for their end users through a modern, developer-friendly approach to security. With their Next-Gen WAF, Bot Management, and DDoS Protections, teams have the tools they need to ensure their applications and APIs perform at their best without sacrificing speed or reliability. Discover how Fastly transforms security into an enabler for digital innovation at <https://www.fastly.com/security>.