



REPORT

State of Application Security 2024



Table of Contents

 [Table of contents](#)

3	Executive Summary	17	Client-side risks
4	Key app security findings		Data snapshot: Third-party scripts and cookie usage
5	Scope and report methodology	18	Business considerations and recommendations
6	Trends in mitigated traffic	19	Shadow API risks
7	Data snapshot: Mitigated traffic over time	20	Business considerations and recommendations
8	Business considerations and recommendations	21	Conclusion
9	Zero-day trends	22	How Cloudflare can help
10	Business considerations and recommendations	23	Learn More
11	DDoS attack trends	24	Appendices
12	Data snapshot: Largest HTTP DDoS attacks		Glossary of key Cloudflare terms
13	Business considerations and recommendations	25	Endnotes
15	Bot traffic trends		
	Data snapshot: Industries with high bot traffic		
16	Business considerations and recommendations		

Executive Summary

Web applications are central to modern life. For governments, they are an important channel to communicate information to the public and provide essential services. For businesses, they serve as a source of revenue, efficiency, and customer insights.

However, the apps and application programming interfaces (APIs) that move critical data, processes, and infrastructure also represent an expanding attack surface. Exploited, unprotected apps can lead to [business disruptions](#), [financial losses](#), and [critical infrastructure collapses](#).

The demand for developers to quickly deliver new features — such as capabilities driven by [large language models \(LLMs\)](#) and [generative AI](#) — magnifies this problem.

Powered by one of the world's largest networks, Cloudflare on average serves over 57 million HTTP requests per second, and blocks 209 billion cyber threats each day as of Q1 2024. The volume, velocity, and variety of this traffic informs the insights explored in this **State of Application Security 2024 report**.

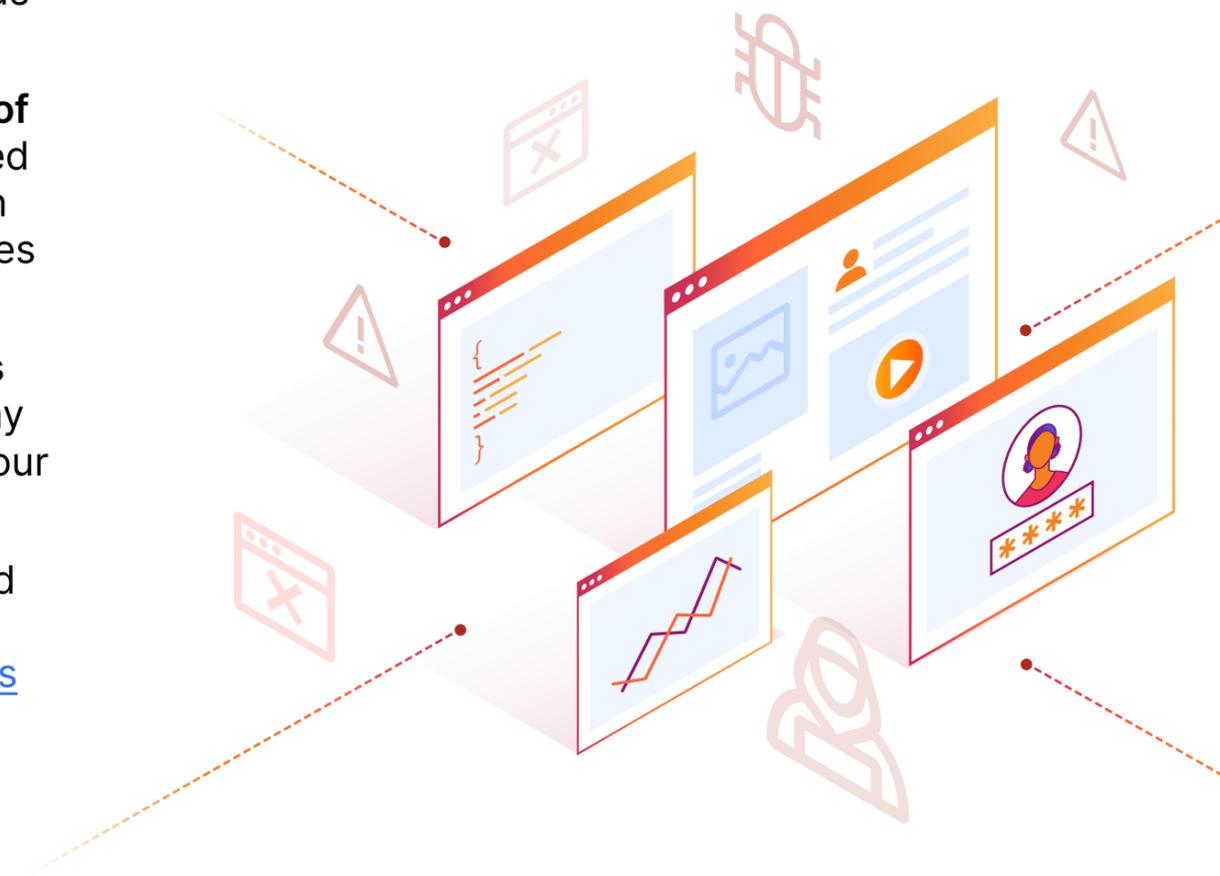
For instance, the increasing speed and volume of [DDoS attacks indicates](#) that **botnets are becoming more common and more efficient at launching DDoS attacks** — the No. 1 attack type against web apps. Is your team equipped to detect and stop traffic from malicious botnets comprising hundreds of thousands — or even millions of machines?

Additionally, certain industries **face a bigger share of bot traffic**. Other industries find themselves targeted by **a high volume of DDoS attacks**. How quickly can you respond to those threats to avoid financial losses and reputation damage?

Cloudflare also found that **enterprise organizations use an average of 47.1 third-party scripts**, as of May 2024. Is your organization inadvertently exposing your end users to supply chain risks?

As new app risks exceed the resources of dedicated app security teams, more organizations recognize the need for a different approach. Gartner® [predicts](#) that, **“By 2027, 30% of cybersecurity functions will redesign application security to be consumed directly by non-cyber experts and owned by application owners.”**

However your organization approaches app security, we hope this report can guide where to prioritize future app security controls — without stifling digital innovation.



Key app security findings

Data collection period: Unless otherwise stated in the endnotes, the time frame evaluated in this report is the 12-month period from April 1, 2023 through March 31, 2024 inclusive.

#1 attack type

Distributed denial-of-service (DDoS) attacks remain one of the most common attack types against web applications, comprising **37.1%** of all app-layer traffic mitigated by Cloudflare.¹

Rapid CVE weaponization

Cloudflare observed an attempted exploitation of a new zero-day vulnerability **just 22 minutes** after its proof-of-concept (PoC) was published.²

Trust in third-party code

Enterprise organizations **use on average 47.1 third-party scripts** — and their web applications make an average of **49.6 outbound connections** to third-party resources.³

93% of bots are potentially malicious

One-third (31.2%) of all traffic stems from **bots**—the **majority (93%) of which are unverified** and potentially malicious.⁴

Outdated approaches to API security

Traditional web application firewall (WAF) rules are most often used to protect API traffic⁵; however, traditional WAF negative security model approaches are **not sufficient** to protect against modern API threats.



Cookie consent risks

Enterprise websites use an **average of 11.5 HTTP cookies** and a median of 5 cookies.⁶ These HTTP cookies may expose end-users to privacy risks that application owners are responsible for monitoring and minimizing.

Scope and report methodology

At a broad level, Cloudflare mitigated 6.8% of all web application traffic during the data collection period.⁷ “Mitigated” traffic is defined as any traffic that is blocked or is served a challenge by Cloudflare (see the Glossary for the full technical definition). The specific threat type and relevant mitigation technique depends on many factors, such as the application’s potential security gaps, the nature of the victim’s business, and the attacker’s goals.

Some examples of attacks on web applications and APIs in 2023-2024 included:

- The [Anonymous Sudan](#) group launched politically motivated DDoS attacks against banks, universities, hospitals, airports, social media platforms, government agencies, and others worldwide.
- Cloudflare observed a record-breaking DDoS attack exploiting a vulnerability in the HTTP/2 protocol, [launched by a botnet](#) of only 20,000 machines that rotated IPs to avoid mitigation.
- T-Mobile [disclosed](#) in early 2023 that it experienced a data breach of 37 million customer accounts via an exploited API.

In other words: the varied nature of such attacks makes web application security a broad discipline that nonetheless requires specialized tools to stop specialized attacks.

To cover such a wide scope, this report is based on aggregated traffic patterns (observed from April 1, 2023 - March 31, 2024) across the Cloudflare global network, including services that:

- Filter HTTP traffic between a web application and the Internet to stop a wide range of real-time attacks using a variety of security measures (*Web Application Firewall*)
- Mitigate DDoS attacks targeting [Domain Name System \(DNS\) servers](#) (*Advanced DDoS Protection*)
- Act as an intermediary to accept, transform, route, and manage all API calls (*API Gateway*)
- Monitor third-party dependencies in a web application that loads in the client browser and exposes the end user to risk (*Page Shield*)
- Identify bot activity, bot reputation, bot origin, and other bot behaviors (*Bot Management*)
- Block users, bots, or applications from over-using or abusing a web property (*Rate Limiting*)



This data and threat intelligence from Cloudflare’s network has been complemented by third-party sources, which readers can access using the inline links.

Trends in mitigated traffic

Compared to the prior 12-month period, **Cloudflare mitigated a higher percentage of application layer traffic and layer 7 (L7) DDoS attacks (6.8% vs. 6%)** between Q2 2023 and Q1 2024.⁸

WAF product mitigations also took over as the No. 1 mitigation technique — a spot that DDoS protections previously held.

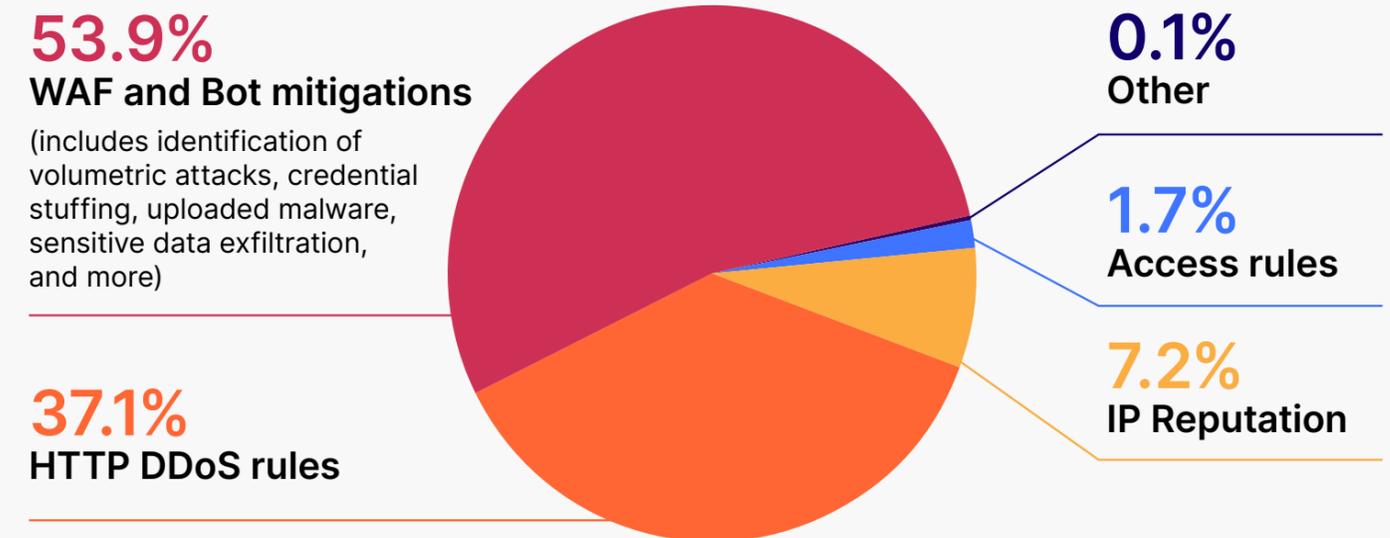
This change in mitigation rankings may be due to more enterprises using WAF rules to block brute-force attacks or credential stuffing and prevent sensitive data from being exfiltrated from apps, or using Cloudflare’s machine learning to block zero-day vulnerability exploit attempts before disclosure.

WAF rules also include **custom rules**, which help enforce organizational policy and perform other custom mitigations.

Some [common use cases for custom rules](#) include:

- Allowing traffic from search engine bots
- Allowing traffic from specific countries only
- Challenging bad bots
- Configuring token authentication
- Requiring a specific cookie

Figure 1: Mitigated traffic by Cloudflare product group⁹

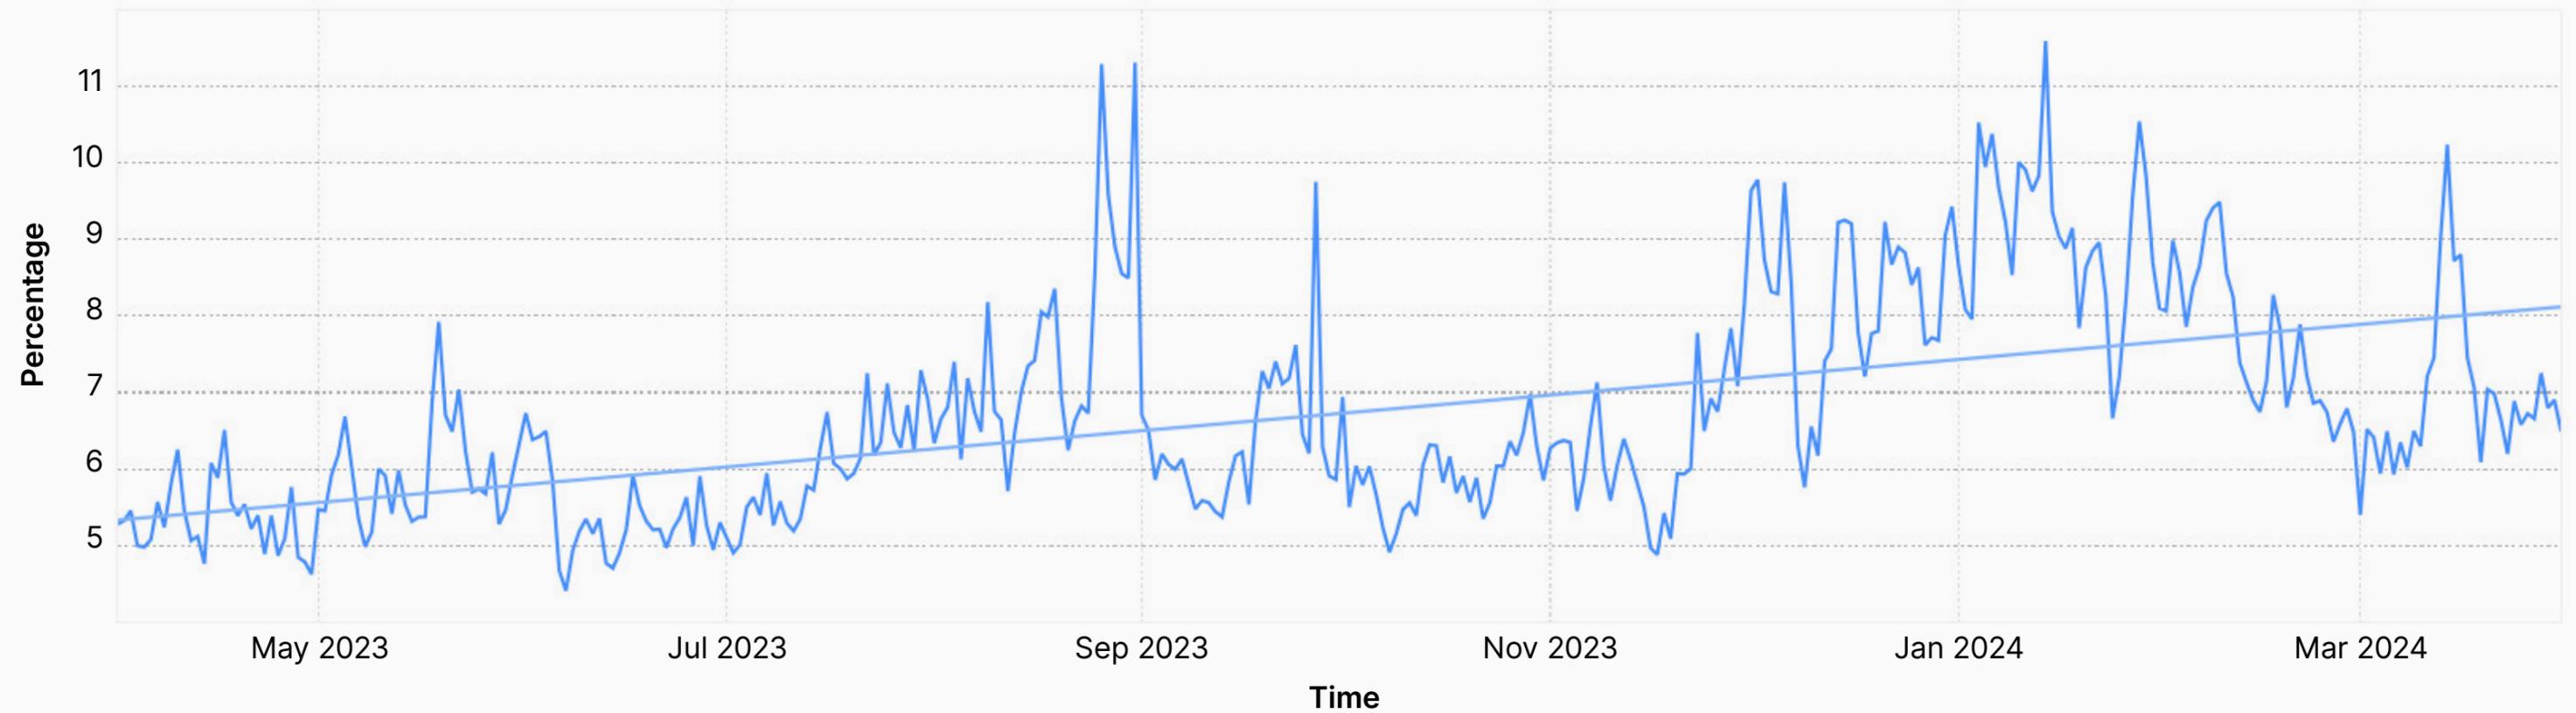


Definitions of these mitigation types can be found in the [glossary section](#).

Data snapshot: Mitigated traffic over time

Cloudflare observed an overall increase in mitigated traffic (and by extension, attacks) in the 12-month period leading up to March 31, 2024. We also saw a spike in attack traffic in January of 2024 this year, and a lower spike during the winter holidays than expected.

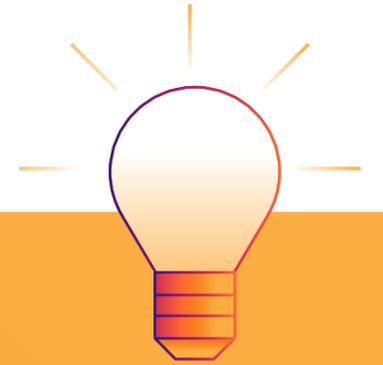
Figure 2: Percent of mitigated HTTP traffic on Cloudflare's global network between Q2 2023 - Q1 2024¹⁰



App traffic shows no signs of decelerating as businesses continue to modernize legacy apps or release new apps:

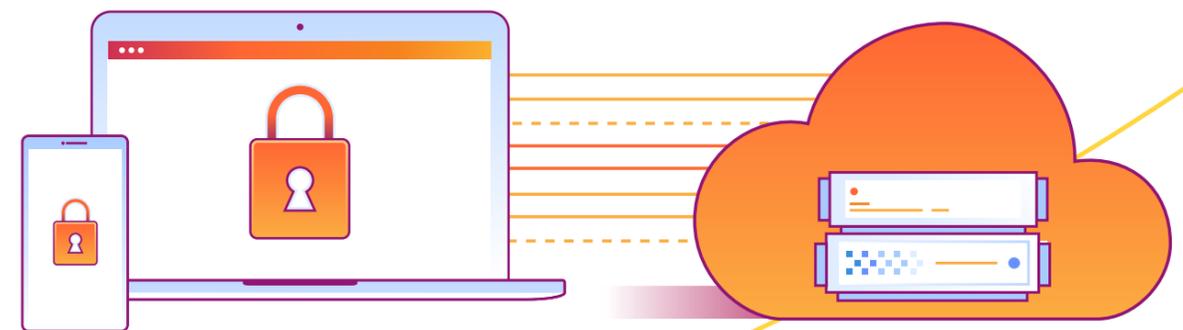
- To improve application performance for globally distributed data and users
- To migrate legacy apps to the cloud, hybrid, or multi-cloud environments
- To augment the user experience with AI-driven insights, recommendations, and information
- To modernize back-office processes and functions
- To build development pipelines from a variety of different tools so that developers can focus on coding

Business-driven app development cannot slow down; therefore, the need to block, challenge, and throttle (i.e., mitigate) malicious or unwanted web app traffic will likely grow.



Recommendations

To help reduce costs associated with scaling infrastructure to serve application growth, organizations should **consider serving application content and mitigating attacks at the edge**. (A group of Cloudflare customers self-reported that they saved an average of about 30% on infrastructure costs by serving and mitigating traffic at the edge).¹¹



Zero-day trends

[Zero-day exploits](#) (also called zero-day threats) are increasing, as is the speed of weaponization of disclosed CVEs.

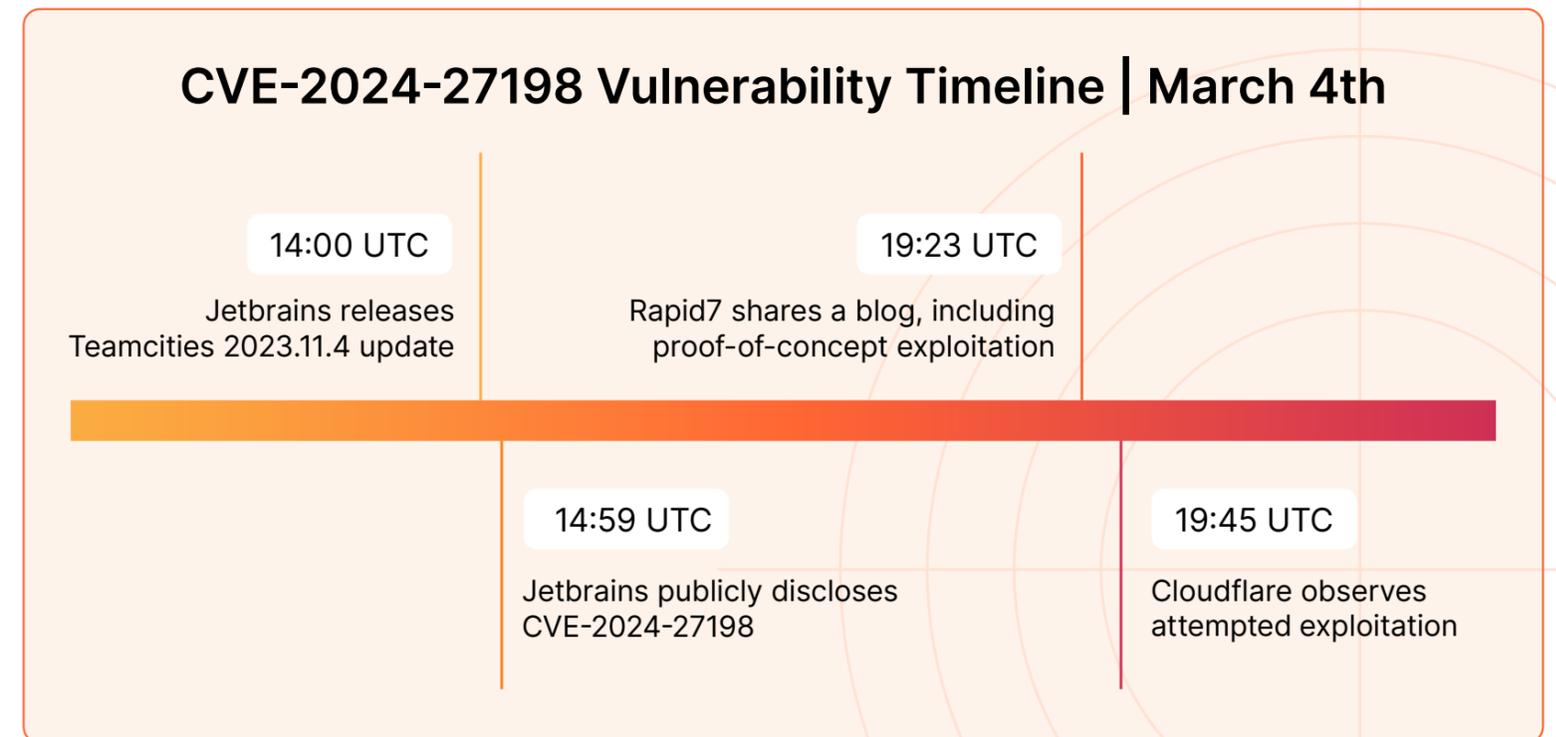
- **97 zero-days** were [exploited in the wild](#) in 2023
- The number of disclosed CVEs between 2022 and 2023 [increased by 15%](#)
- [More than 5,000](#) critical vulnerabilities were disclosed in 2023, yet the mean time to release a patch for a critical severity web application vulnerability is [35 days](#)

Looking at CVE exploitation attempts against customers, Cloudflare mostly observed **scanning activity**, followed by **command injections**, and some **exploitation attempts of vulnerabilities** that had PoCs available online (e.g., Apache, Coldfusion, MobileIron).¹²

This trend in CVE exploitation attempt activity indicates that attackers are going for the easiest targets first, and likely having success in some instances given the continued activity around old vulnerabilities.

The speed of exploitation of disclosed CVEs is often quicker than the speed at which humans can create WAF rules or create and deploy patches to mitigate attacks.

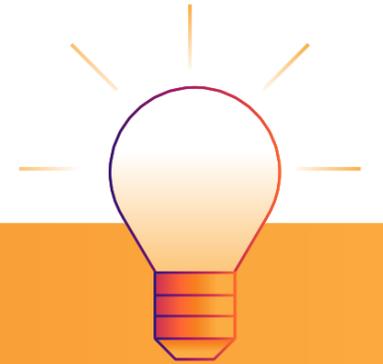
For instance, when Cloudflare observed exploitation attempts of CVE-2024-27198 at 19:45 UTC on March 4, it had taken attackers just **22 minutes after proof-of-concept code was published**.



Business considerations and recommendations

By definition a zero-day is a vulnerability for which there is not a patch in place. A race occurs after disclosure between security professionals trying to defend applications and attackers trying to exploit applications.

The faster novel attack vectors can be detected and mitigated before they cause an issue, the more time internal teams have to patch and remediate the underlying vulnerability. However, sometimes CVE patches are not available for hours (or days or months).



Recommendations

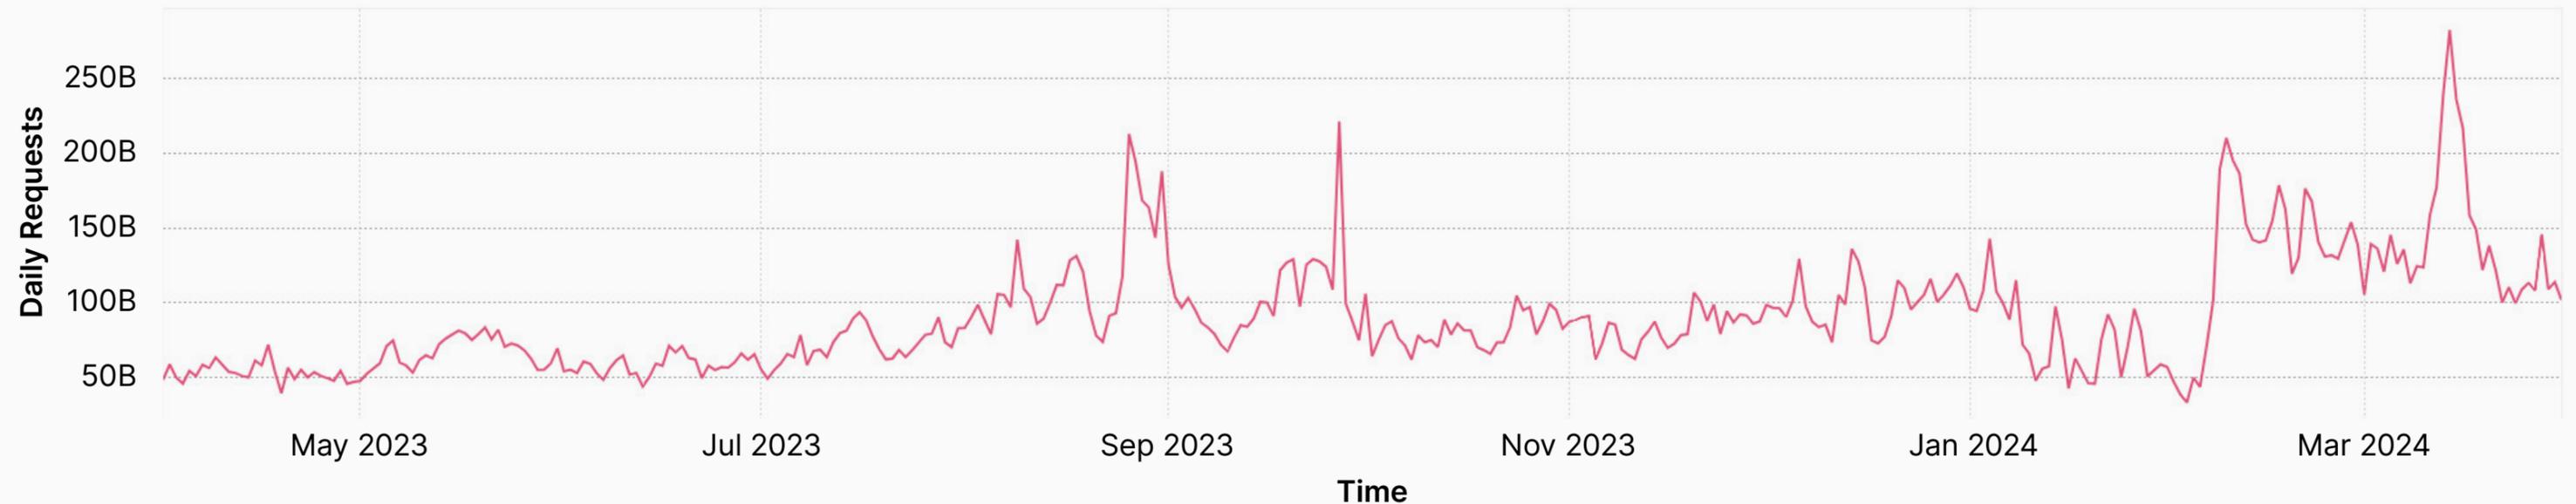
Resource-strapped organizations should **prioritize high-risk and actively-exploited vulnerabilities first, while also using a WAF deployment that provides automatic rule updates** to cover applications that cannot be patched quickly enough.

WAF machine learning (ML) models make it easier to block some zero-day exploits before they are made public and vulnerabilities disclosed.

For example, certain Sitecore CVEs initially disclosed in June 2023 were not initially identified by Cloudflare Managed Rules — but they were correctly detected and classified in ‘zero time’ by our machine learning-based classifiers. Cloudflare also blocked the Ivanti Connect Secure vulnerability before the vulnerability had even been publicly disclosed.

Figure 3: Volume of application DDoS attacks over time¹³

Worldwide - Data ranges from 2023-04-01 to 2024-03-31



DDoS attacks remain the most common attack type against web applications, **with DDoS comprising 37.1% of all mitigated application traffic** (see Figure 1).⁹

We saw a large increase in volumetric attacks in February and March of 2024.¹³ In this first quarter of 2024 alone, Cloudflare's automated defenses mitigated 4.5 million DDoS attacks — an amount equivalent to 32% of all the DDoS attacks Cloudflare mitigated in 2023.

Specifically, application layer HTTP DDoS attacks [increased by 93% YoY](#) and [51% quarter-over-quarter \(QoQ\)](#).¹⁴

As an example, Cloudflare [observed](#) a 466% increase in DDoS attacks on Sweden after its acceptance to the NATO alliance on March 7, 2024. This mirrored the DDoS pattern observed during Finland's NATO acceptance in 2023.¹⁵ **The size of DDoS attacks themselves are also increasing, as illustrated on the next page.**

Data snapshot: Largest HTTP DDoS attacks

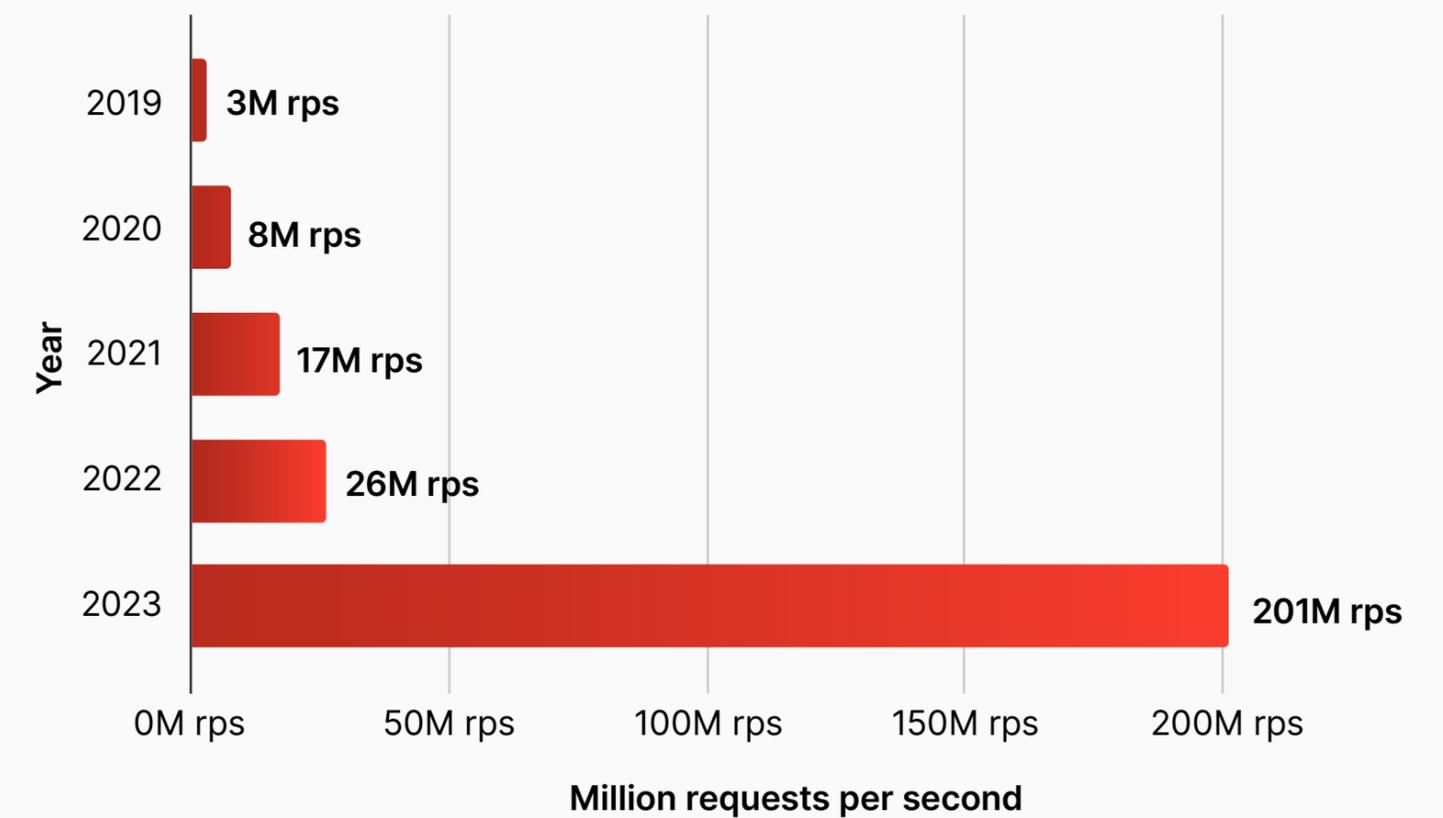
In 2023, Cloudflare mitigated a **hyper-volumetric DDoS attack that peaked at 201 million requests per second (rps)** – three times larger than any previously-observed attack.¹⁶

In the [“HTTP/2 Rapid Reset” attack](#), threat actors exploited a zero-day vulnerability in the HTTP/2 protocol, which is critical to how the Internet and all websites work.

The vulnerability exploit’s potential to incapacitate nearly any server or application supporting HTTP/2 underscored how menacing DDoS vulnerabilities are for unprotected organizations.



Figure 4: Largest HTTP DDoS attacks as seen by Cloudflare, by year



Source: [Cloudflare’s DDoS threat report for 2023 Q4](#)¹⁷

Business considerations and recommendations

HTTP/2 Rapid Reset and other large DDoS attacks illustrate that DDoS attacks are being launched more efficiently by botnets.

For example, cyber crime groups on the dark web offer DDoS-as-a-service for inexpensive prices, even offering “subscribe and save” bundles and support tiers.

Many sites offering DDoS-as-a-service [charge as little as](#) \$10 USD for a DDoS attack that lasts an hour as of 2023, or \$35-170 USD for a full day use of their botnets.

Given how simple it has become for attackers to launch DDoS campaigns, businesses in the following industries should be particularly vigilant about maintaining advanced DDoS protections.

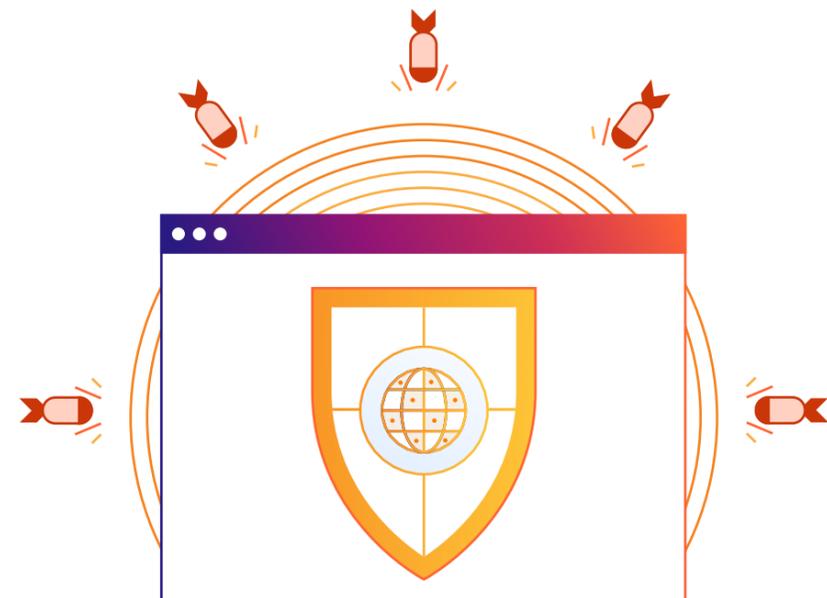


Figure 5: Top industries experiencing L7 DDoS attacks as a share of all Internet traffic¹⁸

1	Gaming and gambling
2	IT and Internet
3	Cryptocurrency
4	Computer Software
5	Marketing and Advertising
6	Telecommunications
7	Retail
8	Adult Entertainment
9	Banking, Financial Services, and Insurance
10	Manufacturing

Business considerations and recommendations (continued)

For DDoS protection provided by a public cloud, a cloud provider typically sits in front of an organization's applications and infrastructure and diverts all traffic to a scrubbing center to be "cleaned." Only legitimate traffic is sent back to the customer.

This motion can be activated either 'on-demand' or 'always-on.' However, there are usually limitations:

- On-demand cloud scrubbing relies on human intervention, adding time to the mitigation response. Providers may also charge per byte of attack traffic, which is costlier over time.
- Many always-on DDoS vendors rely on distant scrubbing centers that can introduce latency and noticeable delays.



Recommendations

To realize the full benefits of cloud-based DDoS defense, look for a scalable, "always-on" service with these capabilities:

- Automatic absorbing of malicious traffic as close as possible to the attack origin, to reduce end-user latency and business downtime
- Unmetered, unlimited DDoS attack mitigations, without charging penalties for spikes in attack traffic
- Centralized autonomous protections against all DDoS attack types



Bot traffic trends

On average, bots comprise one-third (31.2%) of all application traffic processed by Cloudflare.¹⁹ This percentage has stayed relatively consistent (hovering at about 30%) over the past three years.

The term bot traffic may carry a negative connotation, but in reality bot traffic is not necessarily good or bad; it all depends on the purpose of the bots. Some are “good” and perform a needed service — such as customer service chatbots and authorized search engine crawlers. But some bots misuse an online product or service and need to be blocked.

Different application owners may have different criteria for what they deem a “bad” bot. For example, some organizations may want to block a content scraping bot that is being deployed by a competitor to undercut on prices, whereas an organization that does not sell products or services may not be as concerned with content scraping. Known, good bots are classified by Cloudflare as “verified bots.”

However, the vast majority (93%) of bots we identified were unverified bots, and potentially malicious.²⁰

Unverified bots are often created for disruptive and harmful purposes, such as hoarding inventory, launching DDoS attacks, or attempting to takeover an account via brute force or credential stuffing. ([Verified bots](#) are those that are known to be safe, such as search engine crawlers).

Bad bots — if left unchecked — can cause massive problems:

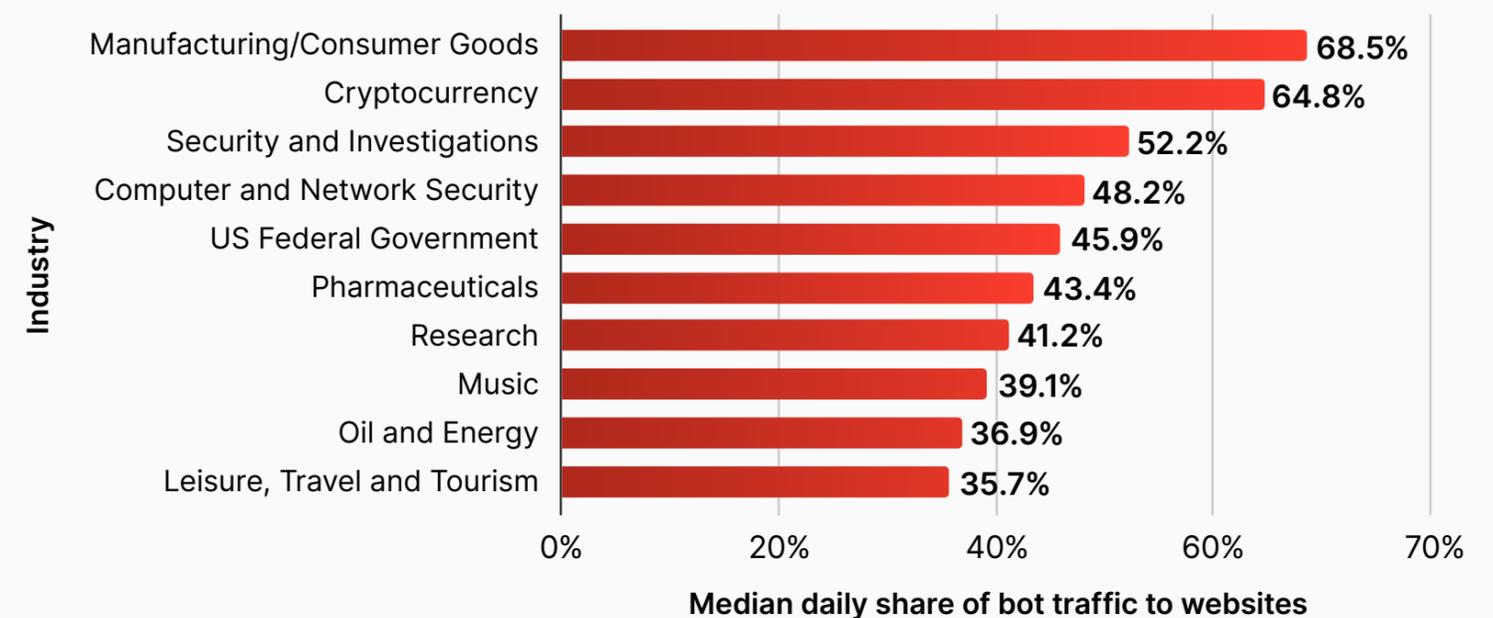
- **Performance impact:** Too much bot traffic can put a heavy load on web servers, slowing or denying service to legitimate users.

- **Business disruption:** Bots can [scrape or download content](#) from a website, rapidly [spread spam content](#), or hoard your online inventory
- **Data theft and account takeover:** Bots can steal credit card data, login credentials, and take over accounts

Data snapshot: Industries with high bot traffic

Attackers leveraging bots focus most on industries that could bring them high financial gains.

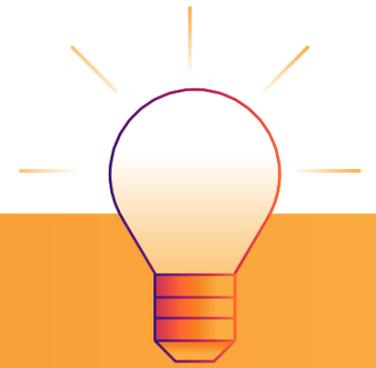
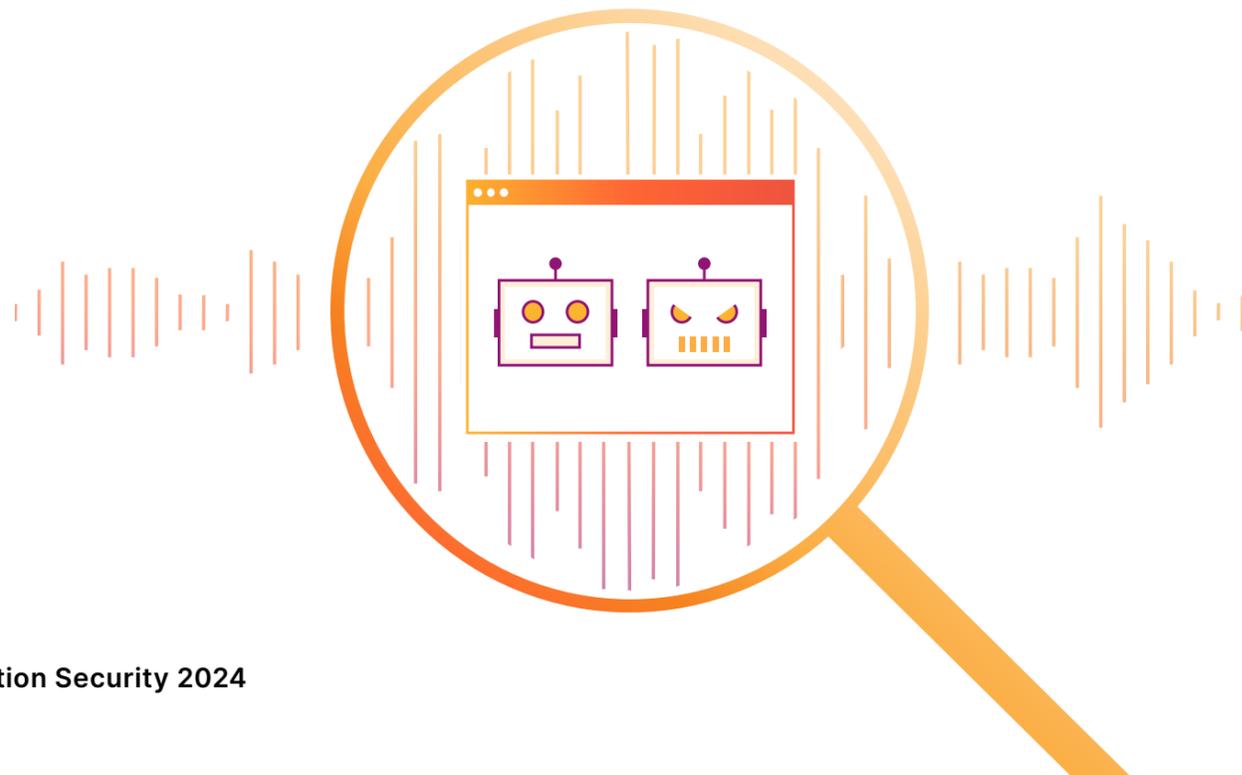
Figure 6: Industries with the highest median daily share of bot traffic²¹



Business considerations and recommendations

As shown in the prior graph, when we analyzed which industries have the biggest bot problem, we found **manufacturing and consumer goods deal with a staggering 68.5% of all traffic** to their websites originating from bots.²¹

Our findings reinforce what has also been [observed](#) by retailers in the consumer goods industry (for example, inventory hoarding bots [buying up shoes](#) or gaming consoles before humans get a chance to put the items in their carts), damaging brand trust. On the other hand, industries that sell fewer physical goods online, such as insurance or hospitality, deal with a bot percentage closer to the Internet average of 31.2%.²¹



Recommendations

If your industry tends to experience more bot traffic, consider **boosting investments in bot management** to preemptively stop credential stuffing, content scraping, content spam, inventory hoarding, and other threats from bad bots.

Look for a bot management service that:

- **Accurately identifies bots** at scale by applying behavioral analysis, machine learning, and fingerprinting to a diverse and vast volume of data
- **Integrates easily** with your other web application security and performance services (e.g., WAF, CDN, DDoS)
- **Allows good bots**, such as those belonging to search engines, to keep reaching your site while preventing malicious traffic

Client-side risks

Most organizations' web apps rely on separate programs or pieces of code from third-party providers (usually in JavaScript). The use of third-party scripts accelerates modern web app development and allows organizations to ship features to market faster, without having to build all new app features in-house.

Data snapshot: Third-party scripts and cookie usage

In fact, Cloudflare's typical enterprise customer uses an average of **47.1 third-party scripts, and a median of 20.0 third-party scripts.**²² The average is much higher than the median due to SaaS providers, who often have thousands of subdomains. **Here are some of the top third-party scripts Cloudflare customers commonly use:**²³

- Google (Tag Manager, Analytics, Ads, Translate, reCAPTCHA, YouTube)
- Meta (Facebook Pixel, Instagram)
- Cloudflare (Web Analytics)
- jsDelivr
- New Relic
- Appcues
- Microsoft (Clarity, Bing, LinkedIn)
- jQuery
- WordPress
- Pinterest
- UNPKG
- TikTok
- Hotjar

While useful, third-party software dependencies are often loaded directly by the end-user's browser (i.e. they are loaded client-side) — placing organizations and their customers at risk given that organizations have no direct control over their security measures. For example, in the retail sector, 18% of all data breaches [originate from Magecart style attacks](#), according to Verizon's 2024 Data Breach Investigations Report.

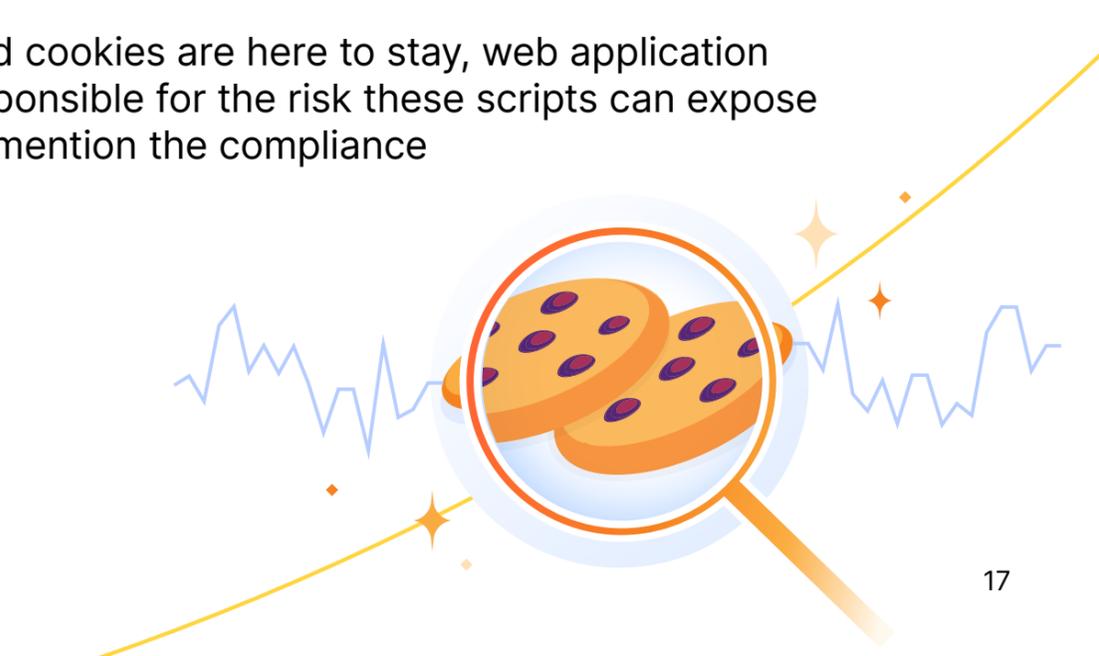
On average, each website has **49.6 connections to JavaScript functions and their destinations, and a median of 15.0.**²⁴ Each of those connections also poses a potential client-side security risk.

Here are some of the top third-party connections Cloudflare customers commonly use:²⁵

- Google (Analytics, Ads)
- Microsoft (Clarity, Bing, LinkedIn)
- Meta (Facebook Pixel)
- Hotjar
- Kaspersky
- Sentry
- Criteo
- tawk.to
- OneTrust
- New Relic
- PayPal

On average, our customers' **websites used 11.5 cookies and a median of five cookies. One organization alone used 131 cookies.**²⁶ Similar to third-party scripts and connections loaded in the browser, cookies also come with both client-side risks and compliance risks. Specifically, cookies can expose website visitors to security risks such as cookie tampering, in which an attacker modifies client-side cookies to perform attacks such as session hijacking in pursuit of account takeover or fraud.

While third-party scripts and cookies are here to stay, web application owners are increasingly responsible for the risk these scripts can expose their end users to — not to mention the compliance and liability implications.



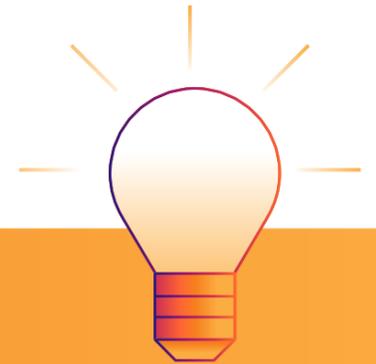
Business considerations and recommendations

Attackers can gain access to [modify the code of JavaScript components](#) used in websites in a variety of ways, such as using stolen account credentials or exploiting [zero-day](#) or unpatched vulnerabilities. Then, they use this privileged access to launch a downstream attack on every website using that JavaScript code.

According to new mandates from [PCI DSS 4.0](#), which will take effect in March 2025, organizations with a payment page are required to monitor third-party script attacks and protect their end-users from browser supply chain attacks.

Cookies also come with client-side and compliance risks (such as cookie tampering, as noted earlier) – if an organization fails to meet user privacy expectations.

The [GDPR's ePrivacy Directive](#), for example, mandates that website owners clearly specify what cookies are being used and for what purposes (and, in some cases, to obtain a user's consent before storing those cookies in the user's browser).



Recommendations

As with client-side scripts, website administrators, developers, or compliance team members do not always know what cookies are being used by their website.

Therefore, **look for a service that automatically neutralizes third-party script risks**, and provides a **full, single dashboard view** of all the first-party cookies being used by your websites.



Shadow API risks

Consumers and end users expect dynamic web and mobile experiences — powered by APIs. For businesses, APIs fuel competitive advantages — greater business intelligence, swifter cloud deployments, integration of new AI capabilities, and more.

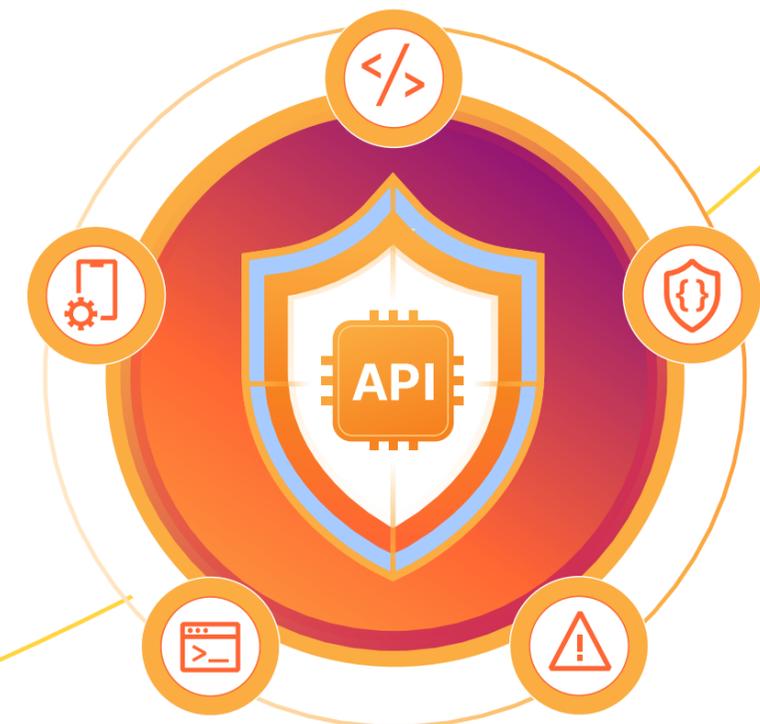
However, APIs — which **now comprise more than half (58%) of the dynamic Internet traffic**²⁷ processed by Cloudflare — introduce new risks by allowing outside parties to access an application.

Yet, for many, API security has fallen behind the fast pace of API deployment: bot operators can directly attack the APIs behind workflows such as account creation, form fills, and payments to steal credentials and more; and AI models' APIs are vulnerable to attacks.

But you cannot protect what you cannot see. **And, many organizations lack accurate API inventories**, even when they believe they can correctly identify API traffic.

Using our proprietary machine learning model that scans not just known API calls, but all HTTP requests (identifying API traffic that may be going unaccounted for), **we found that organizations had 33% more public-facing API endpoints than they knew about.** (This number was the median, and it was calculated by comparing the number of API endpoints detected through machine learning-based discovery vs. customer-provided session identifiers.)²⁸

This suggests that nearly a third of APIs are “shadow APIs” — and may not be properly inventoried and secured



Business considerations and recommendations

Web applications and APIs often work together (for example, an ecommerce website using an API to process payments). However, APIs' unique attributes present a unique attack surface:

	Who interacts with it	Data formats	Request and response structure	Typical threats
 <p>Web apps</p>	Human to system	Flexible (e.g., JavaScript, HTML, CSS)	Flexible, and returns views	DDoS, malicious bots, OWASP Top 10 Web App Risks
 <p>Modern APIs</p>	System to system	Structured and machine-readable (e.g., JSON)	Defined by API schema, and returns only data	Abuse, data exfiltration, malicious bots, OWASP Top 10 API risks

Despite the fact that APIs present different security challenges compared to web apps, we found that 66.6% of API traffic defended by some form of layer 7 security is primarily protected with traditional negative security WAF rules rather than with specialized API rules employing a positive security model.²⁹

Traditional WAF negative security model approaches may be unable to detect all attack traffic directed at APIs, especially API-specific attacks like endpoint enumeration or authentication hijacking. Any WAF used to protect API endpoints should have modern API-specific capabilities that can enforce a positive security model.



Recommendations

As businesses expose more services via APIs, they should augment web app security tools (like WAFs and DDoS), with purpose-built API security and management. Advanced API security, using unsupervised machine learning, helps organizations:

- **Discover shadow APIs:** Constantly scan for every public API in your landscape, even those that are unmanaged or unsecured
- **Prevent data exfiltration:** Stop data leaks by continuously scanning response payloads for sensitive data
- **Create a positive security model:** Protect APIs by only accepting traffic that confirms to your OpenAPI schemas — while blocking malformed requests and HTTP anomalies

The data is clear: the complexity of securing an organization's applications and APIs from new risks continues to grow:

- Application layer HTTP DDoS attacks are increasing by volume and size — and are launched more efficiently by botnets
- The majority of bots are untrusted or unverified, which can negatively impact the security and performance of web apps
- Attackers are weaponizing disclosed CVEs faster; in one example, within just 22 minutes of PoC publication
- Organizations with a higher reliance on third-party scripts and cookies may be at higher risk for software supply chain attacks or privacy and compliance violations

Enterprises often have a disjointed patchwork of legacy and point products for security that make it hard to connect and protect their SaaS apps, web apps, and other IT infrastructure. The IT sprawl makes it easier for attackers to find and exploit vulnerabilities.

The broad nature of web application and API threats requires specialized approaches to stop specialized attacks. However, **a consolidated, best of breed approach helps ensure better security, latency-free connectivity, and business growth.**



How Cloudflare can help

To reduce complexity while protecting more of the growing attack surface, **Cloudflare unifies protections across users, apps, APIs, and networks with a connectivity cloud.**

A connectivity cloud places one unified security network in front of web apps and APIs. This:

- **Stops a wide range of attacks** in real-time using powerful rulesets, exposed credential checks, and other security measures
- **Prevents attackers** from discovering then exploiting IP addresses, configurations, and IT assets
- **Shifts web browsing to the edge** (rather than endpoints), insulating users and devices from web-based threats
- **Detects browser-based attacks**, including client-side attacks that target vulnerable JavaScript dependencies and other third-party scripts

Cloudflare's connectivity cloud scales to protect people, apps, and networks everywhere



~20%
of the web protected
by Cloudflare

320+ cities
in 120+
countries

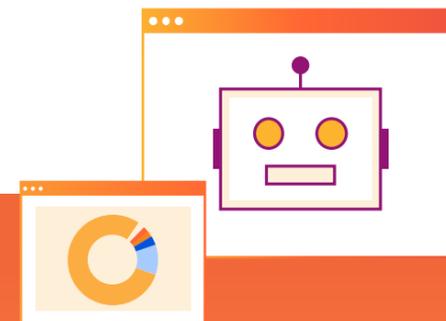
209B+
threats
blocked daily

Learn more

[Table of contents](#)

Built on the backbone of a massive network, our integrated application security portfolio helps organizations take full control of their security posture.

Key services include:



Cloudflare Web Application Firewall (WAF) provides full security visibility, delivers layered protections against OWASP attacks and emerging exploits, detects evasions and new attacks with machine learning, blocks account takeover, detects data loss, and more.

Cloudflare DDoS Web Protection automates intelligent DDoS mitigation from the edge of our global network — mitigating most attacks in three seconds. All plans offer unlimited mitigation of DDoS attacks, with no cost penalty for attack-related traffic spikes.

Cloudflare Bot Management using machine learning, behavioral analysis, and fingerprinting to accurately classify bots. Block credential stuffing, content scraping, inventory hoarding, DDoS, and other malicious bot activity.

Cloudflare API Gateway automatically discovers, validates, and protects your API endpoints. Stop common API attacks, including zero-day exploits, authentication abuse, data loss, DDoS, and other business logic attacks.

Learn more about [Cloudflare's application security and performance solutions](#).

Glossary of key Cloudflare terms

Note: the data in this report is calculated based only on traffic tracked across the Cloudflare network and does not necessarily represent overall HTTP traffic patterns across the Internet.

Custom rules: Allow you to control incoming traffic by filtering requests to a zone. You can perform actions like Block or Managed Challenge on incoming requests according to rules you define.

HTTP DDoS attack rules: A set of pre-configured rules used to match known DDoS attack vectors at layer 7 (application layer) on the Cloudflare global network. The rules match known attack patterns and tools, suspicious patterns, protocol violations, requests causing large amounts of origin errors, excessive traffic hitting the origin/cache, and additional attack vectors at the application layer.

Access rules: Use IP Access rules to allowlist, block, and challenge traffic based on the visitor's IP address, country, or Autonomous System Number (ASN). IP Access rules are commonly used to block or challenge suspected malicious traffic. Another common use of IP Access rules is to allow services that regularly access your site, such as APIs, crawlers, and payment providers.

IP reputation: This threat score measures IP reputation across Cloudflare services. This score is calculated based on Project Honey Pot, external public IP information, as well as internal threat intelligence from our WAF managed rules and DDoS.

Managed rules: Allow you to deploy pre-configured managed rulesets that provide immediate protection against common attacks.

Mitigated traffic: Refers to any eyeball HTTP or HTTPS request that had a “terminating” action applied to it by the Cloudflare platform. This includes the following actions: BLOCK, [CHALLENGE](#), [JS_CHALLENGE](#), and [MANAGED_CHALLENGE](#).

- This does not include requests that had the following actions applied: LOG, SKIP, ALLOW. Starting in 2023, requests that had CONNECTION_CLOSE and FORCE_CONNECTION_CLOSE actions applied by the Cloudflare DDoS mitigation system were also excluded, as these only slow down connection initiation. They accounted for a relatively small percentage of requests.
- Cloudflare improved the calculation regarding the CHALLENGE type actions to ensure that only unsolved challenges are counted as mitigated. A detailed description of actions can be found in the [Cloudflare developer documentation](#).

Rate limiting rules: Allow you to define rate limits for requests matching an expression, and the action to perform when those rate limits are reached.

Uploaded content scanning: When enabled in the Cloudflare WAF, content scanning attempts to detect content objects, such as uploaded files, and scans them for malicious signatures like malware. The scan results, along with additional metadata, are exposed as fields available in WAF custom rules, allowing you to implement fine-grained mitigation rules.

Endnotes

1. Looking at mitigated application traffic between April 1, 2023 - March 31, 2024, we analyzed which application security rules were being used to mitigate traffic. WAF was associated with the most mitigated traffic, but WAF rules block many different types of attacks including volumetric attacks, credential stuffing attacks, malicious content uploads, and more (detected by hundreds of different rules). The second most common ruleset triggered for web applications was the DDoS ruleset, which only identifies DDoS attacks.
2. JetBrains disclosed CVE-2024-27198 on March 4th, 2024 at 14:59. Rapid7 published a proof-of-concept analysis of CVE-2024-2178 several hours later at 19:23 UTC. At 19:45 UTC, Cloudflare observed an attempted exploitation of the vulnerability.
3. We looked at aggregated customer website data pulled from the Page Shield product as of May 2024 with hosts that contained the resource_type = 'script' and resource_type = 'connection' to determine the average number of third-party scripts and connections on each of our customers' hostnames. We eliminated outliers to this dataset, so the number of connections and scripts was found by looking at 99.5% of the dataset.
4. We looked at HTTP traffic to all websites behind the Cloudflare reverse proxy for the report collection period of April 1, 2023-March 31, 2024 and sorted it by human and automated traffic to get the breakdown of bot vs human traffic. To get the breakdown of verified vs unverified bot traffic, we compared the bot traffic against the list that Cloudflare maintains of known "good" bots, also called "verified" bots.
5. To find this data, we looked at application security rules that were triggered for public-facing API endpoints protected by Cloudflare during the collection period of April 1, 2023-March 31, 2024. We then grouped those triggered rules into broader groups corresponding with their products. This helps us get a sense of what tactics attackers are trying most frequently.
6. For the collection period of April 1, 2023-March 31, 2024, we analyzed data from the [URL Scanner project](#), looking at the top 5,000 which received the highest traffic volumes during the data collection period. We chose to analyze the cookies for the top 5,000 domains rather than all URLs behind Cloudflare to show data that is most reflective of the enterprise audience of this report.
7. We analyzed all HTTP requests to applications behind the Cloudflare proxy from April 1, 2023-March 31, 2024 and categorized based on mitigated or non-mitigated (see "Glossary" for definitions of mitigated traffic).
8. We compared the percent of mitigated application traffic from April 1, 2023-March 31, 2024 with data from our report [The State of Application Security in 2023](#).
9. This chart looks at data aggregated during the April 1, 2023-March 31, 2024 period from all applications for which Cloudflare serves as a reverse-proxy and deploys at least one application security rule to identify which security rules are being triggered most frequently. We then grouped those triggered rules into broader groups corresponding with their products. This helps us get a sense of what tactics attackers are trying most frequently.
10. This chart looks at data aggregated during the April 1, 2023-March 31, 2024 period from all applications for which Cloudflare serves as a reverse-proxy and deploys at least one application security rule to identify which security rules are being triggered most frequently. This helps us get a sense of what tactics attackers are trying most frequently.
11. This data is derived from customer feedback during case study interviews, specifically public case studies from [DTLR/Villa](#) and [Open Access College](#), as well as 23 survey responses from a customer return on investment survey run by Cloudflare through the TechValidate software.
12. We looked at attack attempt activity that resulted in a triggered WAF Managed Rule (which is used to stop exploits against common and emerging vulnerabilities) within 30 days of each rule's publication, in order to not overweight Managed Rules released earlier in the year. We examined WAF Managed Rules released between April 1, 2023 - March 31, 2024 and their associated exploit attempt activity.
13. This chart looks at mitigated HTTP traffic over the collection period April 1, 2023-March 31, 2024, zooming in on the volume of mitigations associated with DDoS rules plotted over the year.
14. Source: Cloudflare's [Q1 2024 DDoS Threat Report](#).
15. Source: Cloudflare's [Q1 2024 DDoS Threat Report](#).
16. This data is drawn from [Cloudflare's discovery and analysis](#) of the HTTP/2 vulnerability known as "Rapid Reset" and the subsequent wave of hyper volumetric attacks that resulted.

17. Source: Cloudflare's [Q3 2023 DDoS Threat Report](#).
18. This chart is derived by categorizing HTTP DDoS attacks by industries and then ranking them by which have the biggest share of all DDoS traffic on the Internet during the period April 1, 2023-March 31, 2024.
19. We looked at HTTP traffic to all websites behind the Cloudflare reverse proxy for the report collection period of April 1, 2023-March 31, 2024 and sorted it by human and automated traffic to get the breakdown of bot vs human traffic.
20. To get the breakdown of verified vs unverified bot traffic, we compared bot traffic during the period April 1, 2023 - March 31, 2024 against the list that Cloudflare maintains of known "good" bots, also called "verified" bots.
21. We looked at bot traffic for the period April 1, 2023- March 31, 2024 and categorized it by industry, then compared the share of bot traffic to human traffic for each industry to determine which industries had the highest percentage share of bot traffic.
22. We looked at aggregated customer website data pulled from the Page Shield product as of May 2024 with hosts that contained the resource_type = 'script' and resource_type = 'connection' to determine the average number of third-party scripts and connections on each of our customers' hostnames. We eliminated outliers to this dataset, so the number of connections and scripts was found by looking at 99.5% of the dataset.
23. Using the [Radar Year in Review Report](#) (January 1, 2023 - December 31, 2023) and data from the Cloudflare Page Shield product gathered during the reporting period (April 1, 2023 - March 31, 2024), we collected a list of commonly used third-party scripts that our customers are implementing into their web applications.
24. We looked at aggregated customer website data pulled from the Page Shield product as of May 2024 with hosts that contained the resource_type = 'script' and resource_type = 'connection' to determine the average number of third-party scripts and connections on each of our customers' hostnames. We eliminated outliers to this dataset, so the number of connections and scripts was found by looking at 99.5% of the dataset.
25. Using data from the Cloudflare Page Shield product gathered during May 2024, we collected a list of commonly used third-party connections that our customers are implementing into their web applications.
26. Based on the top 5,000 domains from [Cloudflare's Radar Ranking](#) at the end of 2023. We chose to analyze the cookies for the top 5,000 domains rather than all URLs behind Cloudflare to show data that is most reflective of the enterprise audience of this report. The dataset of Radar's Domain Rankings aims to identify the top most popular domains based on how people use the Internet globally, without tracking individuals' Internet use.
27. Between April 1, 2023 - March 31, 2024, API traffic with successful responses (200 status code) represented a median 58% of Cloudflare's dynamic HTTP traffic. Dynamic content is content that changes based on factors specific to the user, such as time of visit, location, and device.
28. For REST API endpoints, Cloudflare's API discovery tool from the API Gateway product found on median 33% more endpoints through machine learning than we discovered via customer-provided session identifiers across all customers' domains/zones, per account.
29. We examined mitigated traffic to public-facing APIs between April 1, 2023 - March 31, 2024 and checked which products and rulesets were being implemented and triggered most frequently.



© 2024 Cloudflare Inc. All rights reserved.
The Cloudflare logo is a trademark of Cloudflare. All other
company and product names may be trademarks of the
respective companies with which they are associated.

Call: 1 888 99 FLARE
Email: enterprise@cloudflare.com
Visit: cloudflare.com