WHITEPAPER

# The state of application security in 2023

# Content



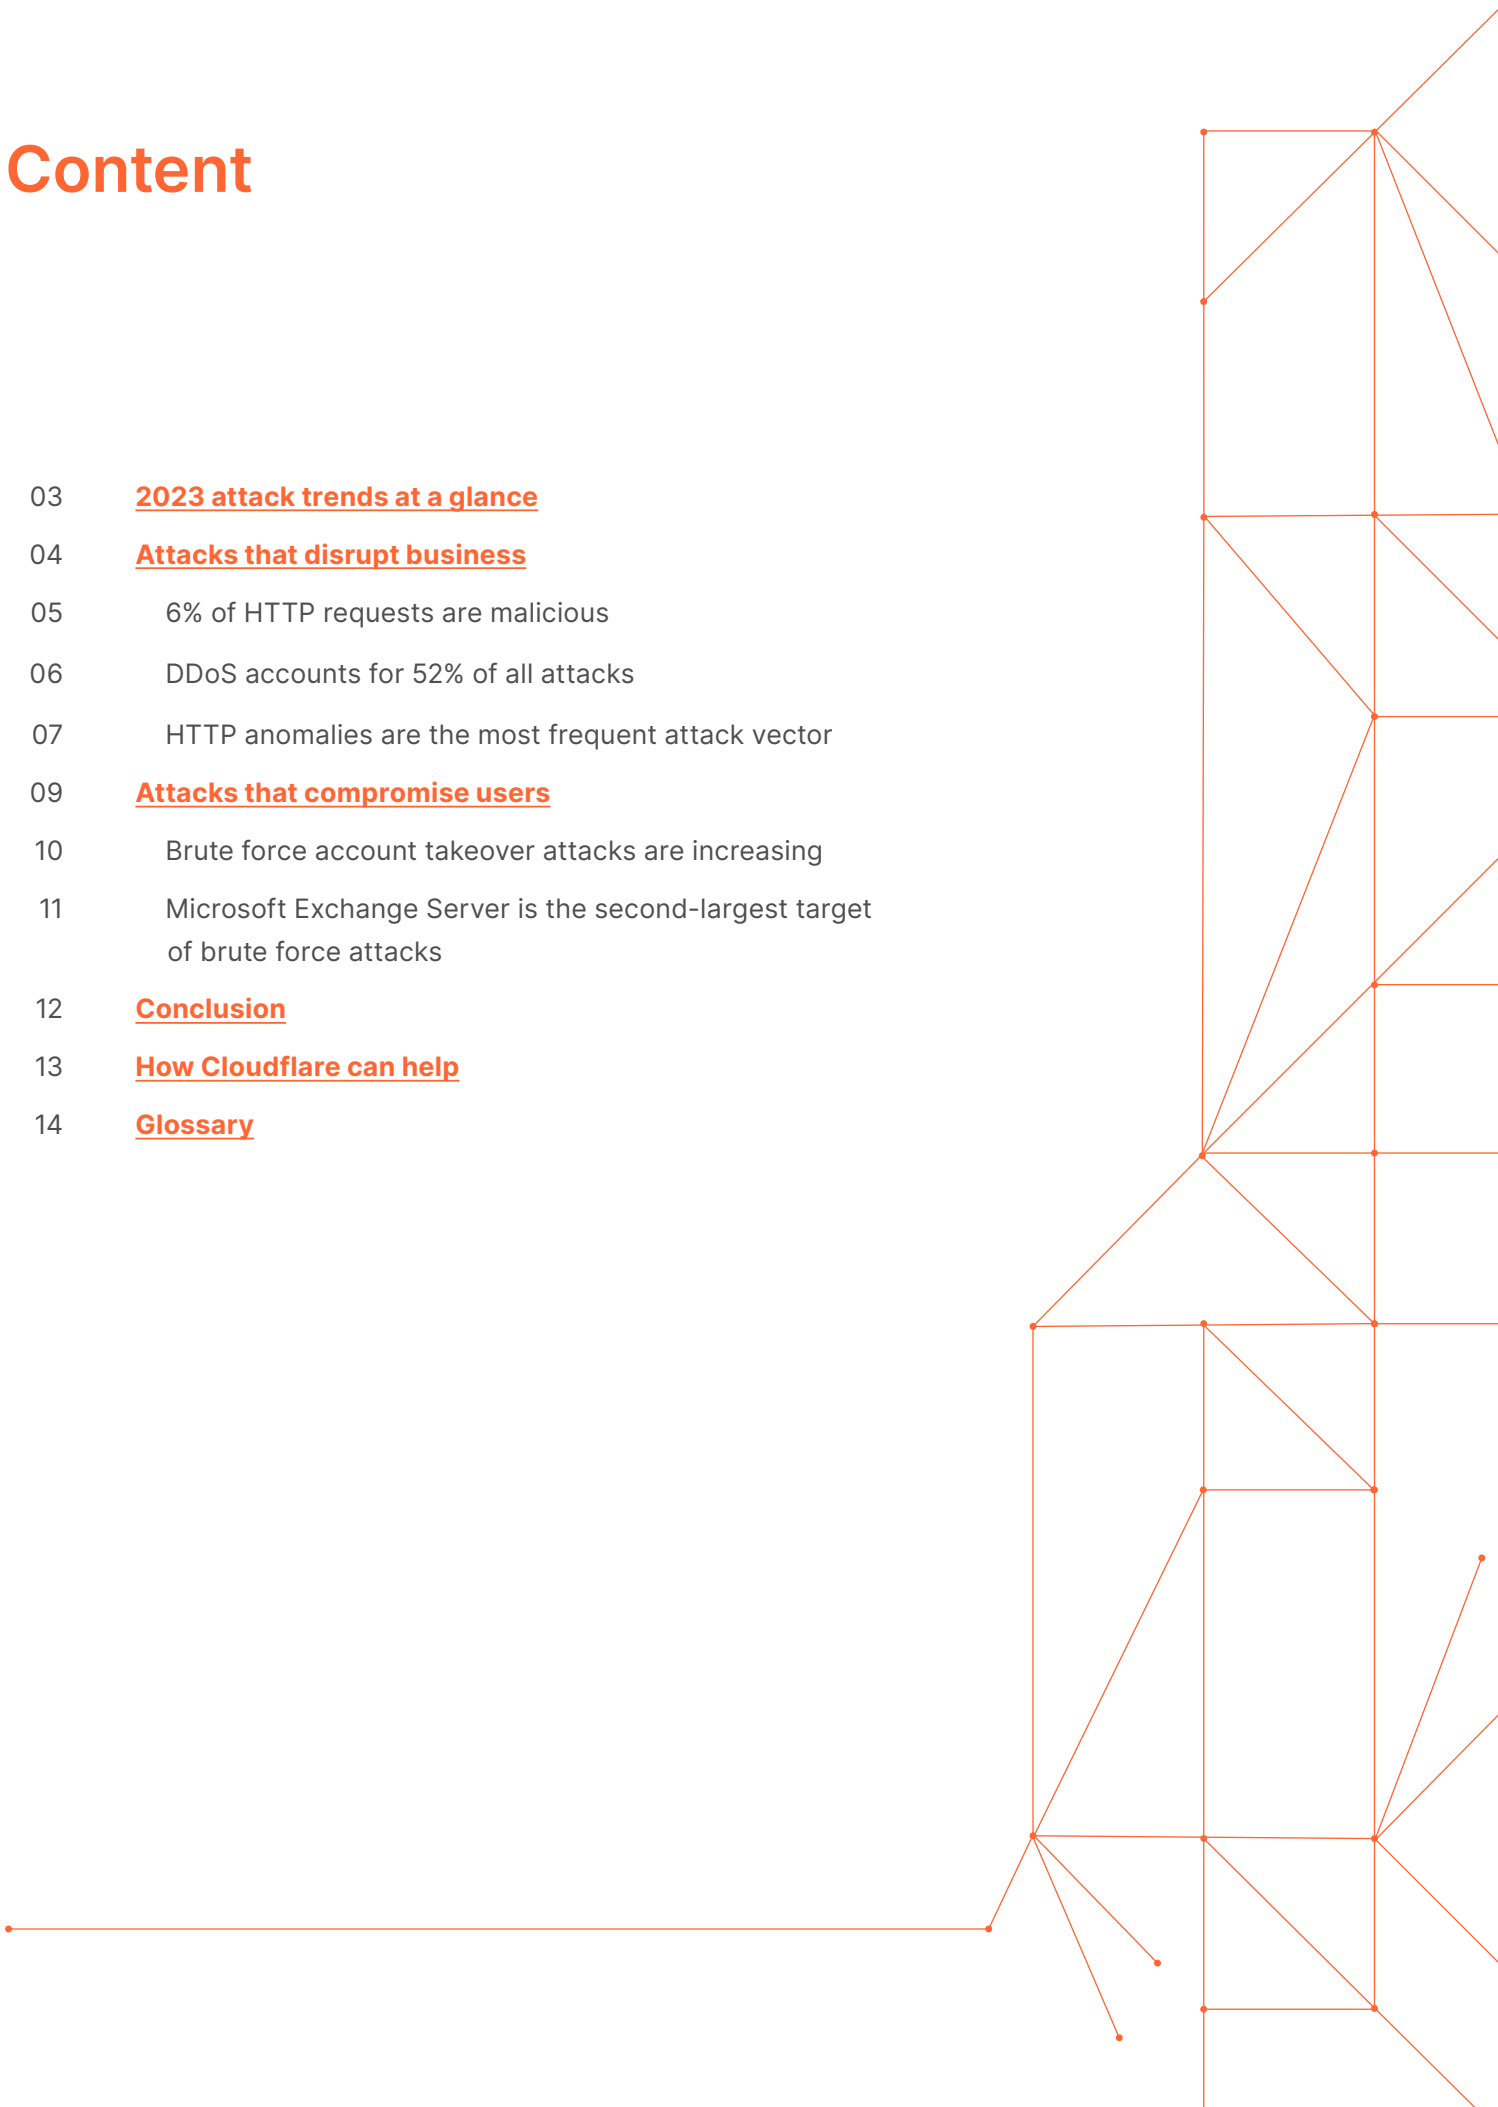

# 2023 attack trends at a glance

It's a line we've all heard before: cyber attacks are larger, more frequent, and more sophisticated than ever. And in 2023, the numbers continue to back up each of these points:

- **Attacks are larger.** In February, **Cloudflare mitigated a 71 million request-per-second HTTP DDoS attack** — the largest-known attack of its kind to date, more than 54% higher than the previous reported record of 46 million RPS in June 2022. To put this into perspective, Google fields approximately 100,000 requests per second across all platforms worldwide, making the attack roughly 140x Google's total traffic.
- **Attacks are more frequent.** **Application-layer attacks have spiked by as much as 80% in 2023.** One reason behind this jump: attackers are leveraging existing Internet infrastructure to amplify their attacks, making them both easier and cheaper to carry out.
- **Attacks are more sophisticated.** As organizations continue to refine their security strategies, attackers evolve their tactics to get around even the most robust defenses. One way of doing that is via brute force attempts, which can aid attackers in gaining access to user accounts and sensitive data. From 2022 to 2023, **Cloudflare observed matches for HTTP requests with leaked credentials at a rate of 12,000+ per minute.**

In the report below, we will dive into the current and emerging attack trends aimed at applications and APIs in 2023. These observations are powered by the Cloudflare global network, which handles an average of 45+ million HTTP RPS — giving us an unprecedented view into traffic patterns, attacker behaviors, and more.

## What we mean by 'HTTP' traffic

Throughout this report, we frequently refer to 'HTTP' traffic. It should be noted that for the purposes of our research, this includes both HTTP and HTTPS traffic.

# Attacks that disrupt business

Most of the attacks observed over the past year have been aimed at businesses and organizations around the globe — with an intent to damage brand reputation, steal sensitive data, or worse. One of the most notorious groups behind these attacks is Killnet.

Killnet, a pro-Russian hacktivist group, commits acts of cyberwarfare via botnet-powered DDoS attacks. Since late 2022, members of the group have threatened to compromise US, UK, and other European healthcare services and government websites. Often, these attacks were launched via publicly-available DDoS scripts and IP stressers, and resulted in several frequent, albeit relatively short, outages.

Killnet is far from the only attacker group threatening organizations and governments today. But their recent behavior is representative of some alarming trends, as attackers continue to launch both multi-terabit, network-level DDoS attacks and waves of smaller, frequent application-level attacks designed to wear down the defenses and consume resources of traditional security appliances.

The effects of these attacks can be devastating, expensive to remediate, and have far-reaching consequences for data compliance. In 2022, the average data breach cost organizations around $4.35 million USD, with some estimates anticipating a $5 million USD price tag in 2023. The cost of remediation may rise even further for businesses who are beholden to strict data privacy regulations, such as the General Data Protection Regulation (GDPR), the California Privacy Rights Act (CPRA), and other laws that enforce hefty fines for compromising user data.

In addition to large-scale DDoS attacks and data breaches, attackers continue to carry out smaller attacks designed to slow down applications, infiltrate organizations, or harm end users. SQL injection (SQLi) and directory traversal attacks, for instance, may be launched in order to exfiltrate sensitive data or create entry points into an organization, while cross-site scripting (XSS) is often used to spread malware or steal information from unsuspecting users.

## Keep reading to find out why...

- 6% of HTTP requests are malicious
- DDoS accounts for 52% of all application attacks
- HTTP anomalies are the most frequent attack vector

## 6% of HTTP requests are malicious

So far in 2023, approximately 6% of total HTTP requests proxied by the Cloudflare network have been mitigated by our security products. The Cloudflare network processes approximately 45+ million HTTP requests per second, and about 2.7 million of these are malicious.

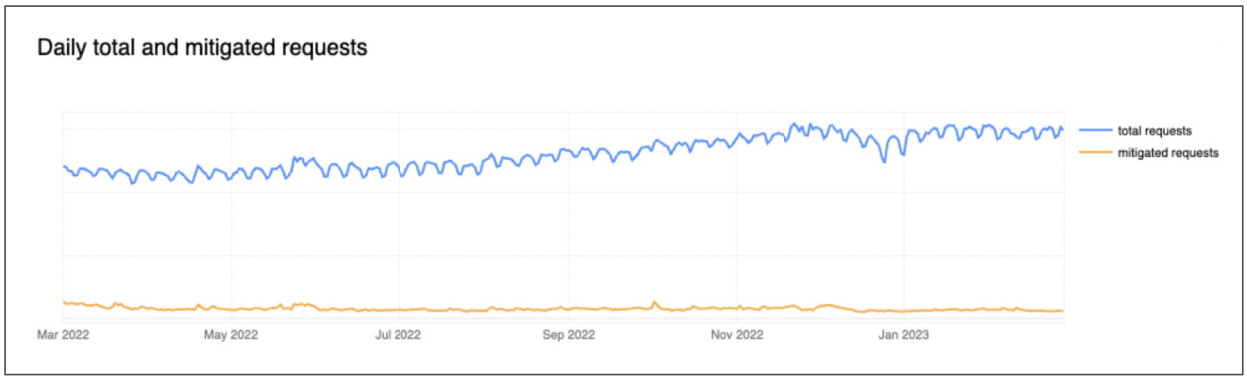


Although the percentage of mitigated traffic has decreased over time, the total mitigated request volume has been relatively stable, indicating an increase in clean traffic globally — rather than an absolute decrease in malicious traffic. In other words: we aren't seeing fewer attacks.

Large spikes visible in the chart below, such as those seen in June and October, often correlated with large DDoS attacks mitigated by Cloudflare.

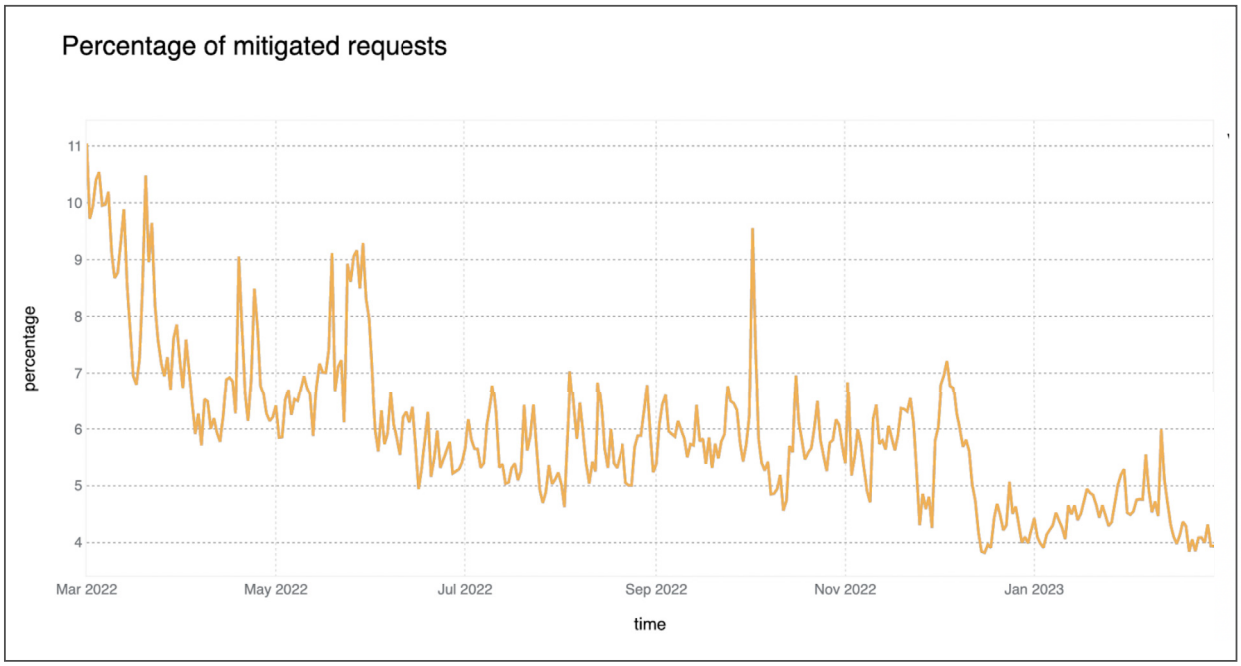


81% of mitigated HTTP requests were outright BLOCKed, with mitigations for the remaining set split across the various CHALLENGE type actions — which, depending on the characteristic of the request, could be anything from a background JavaScript challenge to asking users to click a button.

## DDoS accounts for 52% of all application attacks

DDoS mitigation and web application firewalls (WAF) are two of the strongest defenses against cyber attacks.

Over the last year, DDoS has comprised 52% of all application-layer attacks mitigated by Cloudflare. By contrast, the Cloudflare WAF has been used to mitigate approximately 41% of all application-layer attacks.

The sizable increase in WAF mitigation can be partially attributed to notable advances in the Cloudflare WAF technology that enable it to detect and block a wider range of attacks.

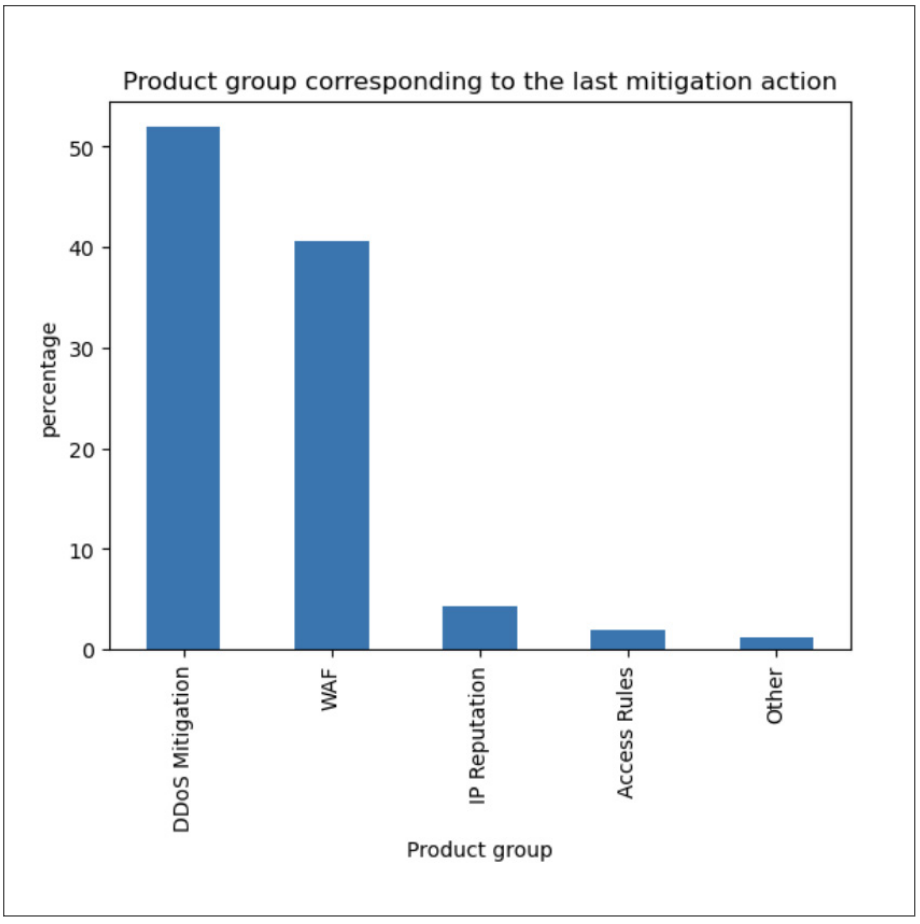


## How Cloudflare mitigated a 71 million request-per-second DDoS attack

In February 2023, Cloudflare reported the largest-known HTTP DDoS attack* — 54% higher than the previous reported attack of 46 million RPS in June 2022.

The attack originated from numerous cloud providers, and targeted popular gaming providers, cryptocurrency companies, and cloud computing platforms, among others.

In response, Cloudflare worked with cloud providers to crack down on the botnet and created a free botnet threat feed for affected providers that own their own autonomous systems.

*This particular attack is not visible in current graphs because the attacks shown are aggregated at a daily level, and this attack only lasted for ~5 minutes.

Another factor that contributed to this increase is a subsequent rise in WAF Custom Rules (formerly “Firewall Rules”) usage. More frequently, customers are augmenting rulesets written by Cloudflare with their own Custom Rules to mitigate malicious traffic that fits within their unique business use cases.

In the following chart, the orange line (firewallrules) reflect this gradual increase over time, while the blue line (l7ddos) trends lower.

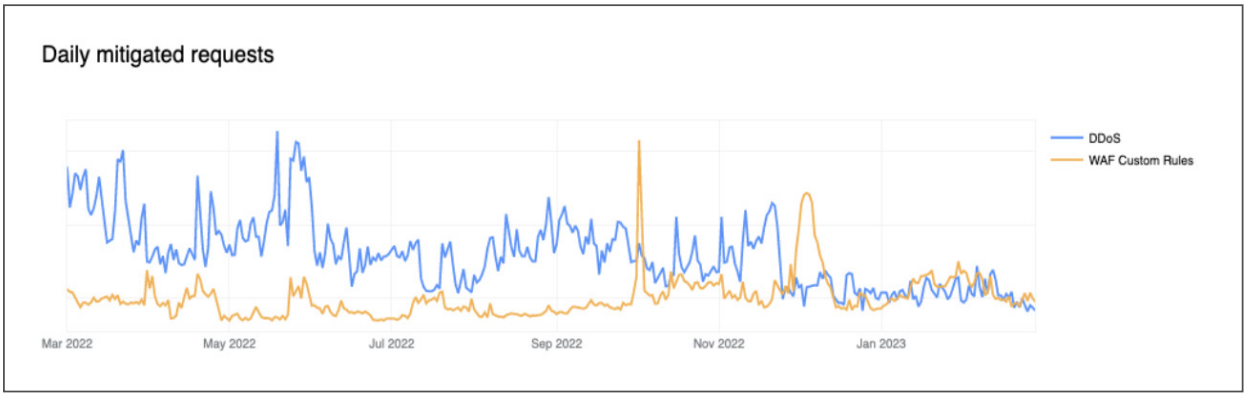

## HTTP anomalies are the most frequent attack vector

Over the last year, HTTP anomalies have become the most frequent application-layer attack vector mitigated by the Cloudflare WAF.

Broadly speaking, HTTP anomalies represent a fairly large category of non-standard HTTP requests that can slow down or harm applications and APIs. These requests may include any number of formatting errors, from missing user agents to the use of non-standard ports. Often, they also reflect attack attempts; as such, it is recommended that customers use custom rules to protect any potentially malicious, non-standard requests.

It is worth noting that the API attack vectors tracked by Cloudflare show considerable variation when compared to global HTTP attack vectors. For example, the following graph displays a notable spike in file inclusion attack attempts around January 2023. Other popular API attack vectors include directory traversals (16%), SQLi (14%), and software-specific attacks (10%).

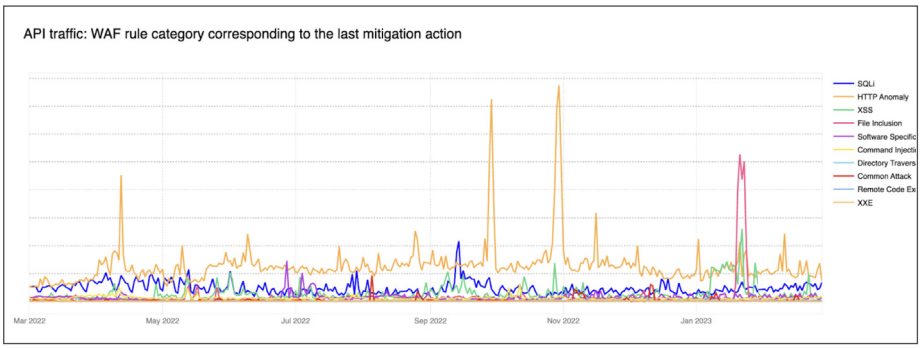


These fluctuations are reflective of the wide variety of traffic handled by the Cloudflare network. A lot of the malicious traffic we mitigate extends beyond standard cross-site scripting (XSS) and SQLi attacks, as attackers are constantly trying new (albeit not always sophisticated) tactics to break applications and APIs.

## How Cloudflare helps block HTTP anomalies

The majority of HTTP requests blocked by Cloudflare Managed WAF Rules contain HTTP anomalies.

These rules include the following:

- Missing user agent: Block requests without a User-Agent header
- Not GET, POST, or HEAD method: Block any other method of request — including those from vulnerability scanners
- Missing referer: Block requests without a REFERER header
- Non-standard port: Block requests to access non-standard ports (such as 80 and 443)
- Invalid UTF-8 encoding: Block requests that contain “special” characters not allowed in UTF-8 encoding

Cloudflare API Gateway also blocks HTTP anomalies in API traffic using positive security model enforcement.

# Attacks that compromise users

Protecting against volumetric DDoS attacks, malicious bots, and other automated attacks requires a robust defense — including cloud-based DDoS mitigation services, web application firewalls, and bot management capabilities.

But what happens when these attacks become even more sophisticated?

Account takeover, brute force attempts, and other advanced tactics can have a devastating effect on organizations' revenue and reputation. Once attackers infiltrate a legitimate user's account, they may steal sensitive data, execute malicious code across a network, or launch even larger attacks.

Although the targeted nature of these attacks makes them more sophisticated than run-of-the-mill DDoS threats, they are often made possible by weak authentication measures, zero-day exploits, and basic human error.

Over the last year, Cloudflare has observed a noticeable uptick in some of these attacks; notably, brute force account takeover attempts. Brute force attacks work by spamming login endpoints with potential username and password combinations until access is obtained. Typically, attackers will test different combinations based on personal information that is easily searchable or made publicly available online. Once they find a combination that works, attackers may take over legitimate accounts and carry out any number of malicious actions.

Even more alarming: attackers frequently use leaked credential data — usernames and passwords that were exposed in previous data breaches — to commit these attacks. For example, if an ecommerce site is breached, an attacker can use the list of exposed passwords to gain entry to customers' bank or mobile payment accounts, an effective tactic considering users will unfortunately reuse passwords across multiple sites.

Preventing these attacks requires a multi-pronged strategy, in which both organizations and individual users should practice good password hygiene, update passwords frequently, use multi-factor authentication (MFA) to secure their accounts, and block potentially malicious requests that use known exposed credentials.

## Keep reading to find out why...

- Brute force account takeover attacks are increasing
- Microsoft Exchange Server is the second-largest target of brute force attacks

## Brute force account takeover attacks are increasing

Beginning in the latter half of 2022, more attackers have targeted login endpoints in an attempt to gain access to user accounts. During large brute force attacks, Cloudflare observed matches for HTTP requests with leaked credentials at a rate of 12,000+ per minute.

These requests are designed to wear down organizations' authentication systems — and with leaked credentials already at their fingertips, attackers may have a much higher success rate.

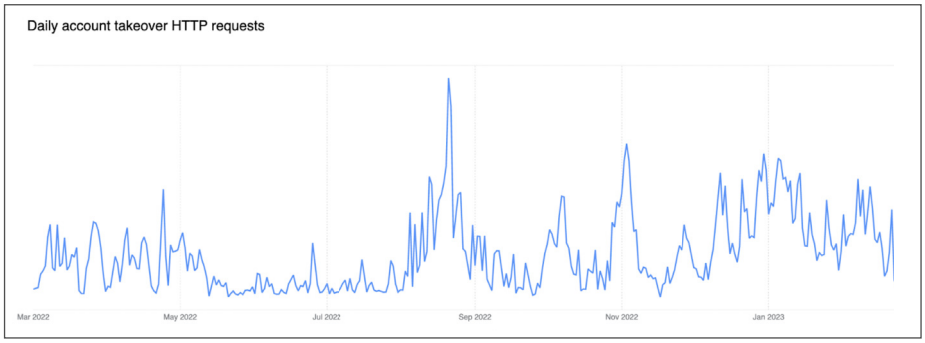


## How Cloudflare defends against account takeover

Cloudflare's exposed credential check feature helps prevent fraudulent account access attempts.

It uses rules that match authentication requests for the following systems:

- Drupal
- Ghost
- Joomla
- Magento
- Plone
- WordPress
- Microsoft Exchange
- Generic rules matching common authentication endpoint formats

Tracking these requests helps Cloudflare identify and block requests using known leaked username/password combinations.

## Microsoft Exchange Server is the second-largest target of brute force attacks

Most applications targeted by brute force takeover attacks tend to be high-value assets. Popular applications — such as WordPress — attract a high number of users and, as such, face an increased risk of data breach and account takeover attempts.

From mid-2022 through the first several months of 2023, Cloudflare has observed a high number of brute force attacks aimed at Microsoft Exchange Server. This calculation was made after observing rule matches from supported systems (Drupal, Ghost, Joomla, Magento, and others). While generic signatures had the most matches for brute force account takeover traffic, Microsoft Exchange accounted for the second-most frequent match.

This kind of attack is particularly dangerous, as attackers used zero-day exploits (i.e. a vulnerability for which there is no known patch) to remotely access targeted networks and execute malicious code. It also is not a new trend, as previous zero-day vulnerabilities identified in Microsoft Exchange software were used to commit cryptojacking, ransomware threats, and other attacks.

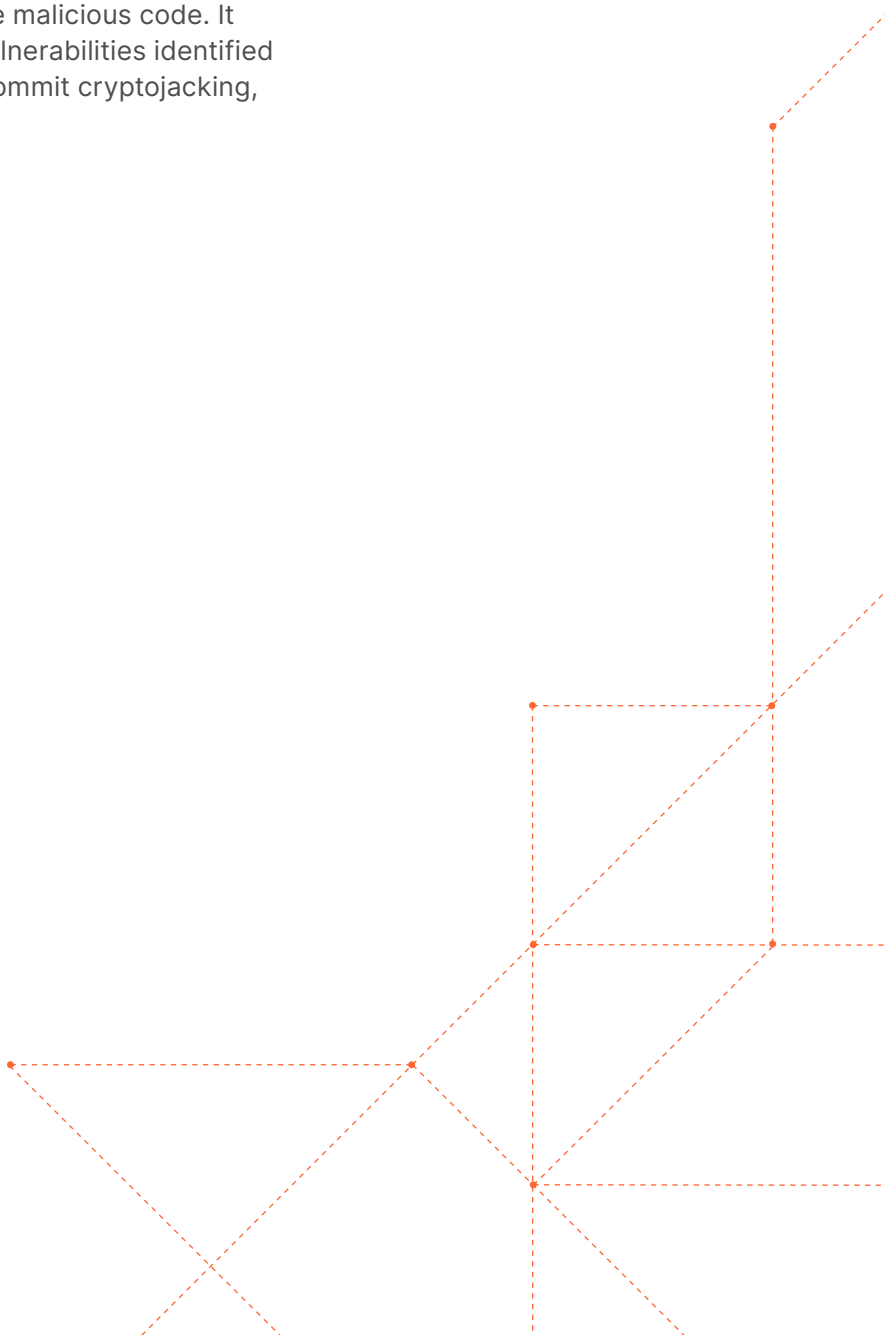

# Conclusion

The threat landscape is constantly changing.

As such, it is critical that any organization operating online continues to invest in threat detection and mitigation technologies so that they can ensure their applications — and, more importantly, their end users' data — remains safe.

The insights provided in this report are a brief sampling of the traffic patterns and attack trends observed on the Cloudflare global network over the last calendar year. Not only does this data help keep organizations informed of emerging and potential threats, but it allows Cloudflare to further refine and strengthen our application security portfolio — so we can continue to help build a better Internet for everyone.

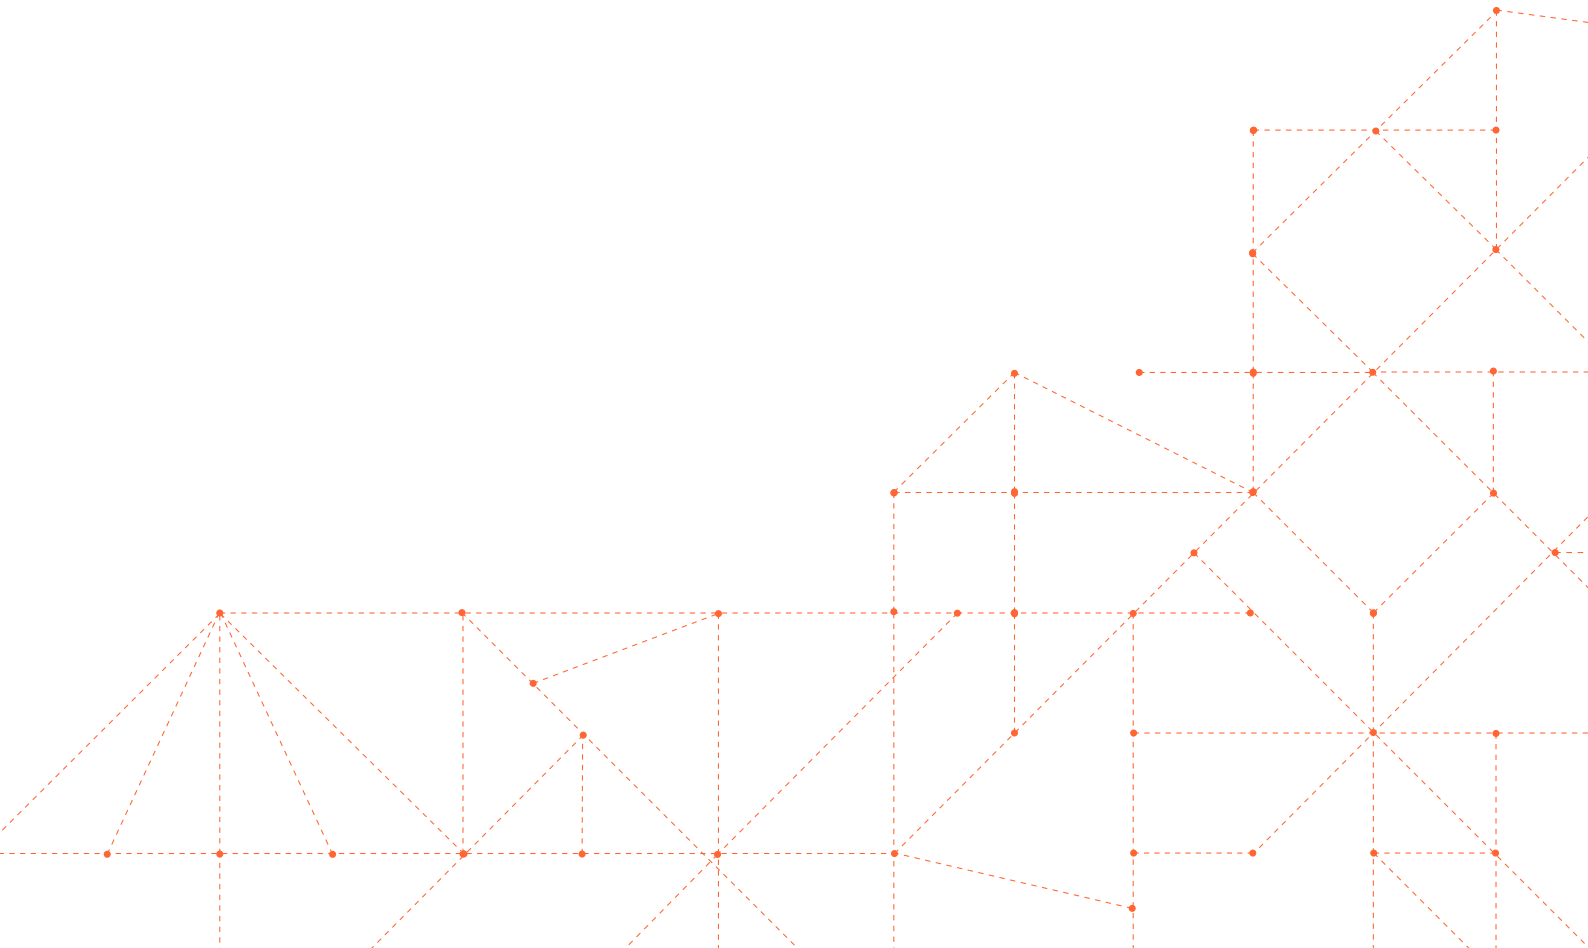

# How Cloudflare can help

Cloudflare operates a global network that spans locations in 285+ cities and over 100 countries. This allows us to observe traffic patterns on a massive scale and better protect our customers from a wide range of both automated and targeted attacks.


Built on the backbone of this network, our integrated application security portfolio helps organizations take full control of their security posture. Here are some of the improvements we are making:

**Cloudflare Web Application Firewall** offers tailored protections to defend against zero-day exploits, Top 10 OWASP attacks, and other threats. It also includes an exposed credential check feature that notifies customers (via an HTTP request header) of leaked username/password combinations. This feature has proven highly effective at detecting botnet-powered brute force attacks.

**Cloudflare DDoS Protection** is fueled by the intelligence of our global network, which allows us to identify and block anomalous traffic from every server we host. All Cloudflare plans offer unlimited, unmetered mitigation of DDoS attacks, regardless of size — with no cost penalty for attack-related traffic spikes.

**Cloudflare Bot Management** continues to refine its ability to track and mitigate malicious bot behavior. In 2023, new features — including configurable heuristics, hardened JavaScript detections, automatic machine learning model updates, and Turnstile (Cloudflare's free CAPTCHA replacement) — have helped improve the classification of human and bot traffic.

**Cloudflare API Gateway** discovers, manages, and protects API endpoints from abusive attacks. With API discovery, organizations can identify, catalog, and monitor potential shadow APIs, while API abuse and sensitive data detection capabilities help pinpoint anomalies and prevent data leaks and other attacks.

# Glossary

- **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.
- **API traffic:** Refers to any HTTP request with a response content type of XML or JSON. Where the response content type is not available (e.g. for mitigated requests), the equivalent Accept content type (specified by the user agent) is used instead. In the latter case, API traffic will not be fully accounted for, but still provides an adequate representation for the purposes of gaining insight into traffic patterns.
- **Application layer:** Refers to attacks that occur at layer 7 of the OSI model.

Unless otherwise stated, the time frame evaluated in this report is the 12-month period from March 2022 through February 2023 inclusive.

Finally, please note that the data is calculated based only on traffic tracked across the Cloudflare network and does not necessarily represent overall HTTP traffic patterns across the Internet.



1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

REV:BDES/4604.2023MAY30