CLOUDFLARE

2026 Cloudflare report on cyberattacks against civil society

Key findings

Distributed-denial-of-service (DDoS) attacks were the most common cyber threat against civil society organizations protected under Project Galileo, accounting for 81.7% of all malicious traffic. Their defining feature was duration. While most DDoS attacks Cloudflare mitigated for its customers were over within minutes, nearly every one of the largest attacks against civil society lasted longer, with some spanning into days and weeks. The Iraq-based digital rights organization Tech4Peace experienced an eight-day long DDoS attack that featured 2.6 billion malicious traffic requests.

On average, Cloudflare blocked a malicious request probing a media organization every seven seconds. In general, civil society organizations faced attempts to exploit security vulnerabilities in websites at a rate more than seven times higher than other Cloudflare customers. Media organizations, including journalists, were the most frequently targeted, receiving 40.5% of attacks, despite making up only 22.7% of the underlying population.

Journalists operating in exile faced a rate of malicious traffic that was nearly four times higher than journalism organizations overall. Attacks were concentrated against a few targets. In December 2025, elTOQUE, a Cuban media outlet operating in exile, faced a DDoS attack that the organization believes was an intentional effort to limit access to a tracker comparing the Cuban peso with foreign currencies.

Nearly 10 percent of all emails Cloudflare processed for civil society included potential phishing material. Compared to other Cloudflare customers, civil society faced a higher concentration of malicious emails intended to gain unauthorized access. Traditional authentication protocols alone left civil society organizations exposed. Nearly one in three emails that contained malicious content bypassed standard authentication methods but were identified by more sophisticated phishing detection tools provided by Cloudflare.

Cloudflare identified 183 Internet disruptions across its global network, 85 of which public reporting has attributed to government action. The restrictions coincided with periods of elections, protests, and student exams. In countries like Iran and Uganda, civil society organizations reported that shutdowns disrupted their ability to reach affected communities, document abuses, and share independent information

Introduction

Cloudflare provides free cybersecurity services to more than 3,400 domains belonging to organizations in 120 countries through Project Galileo. Because of the number of organizations in the program and their geographic representation, Cloudflare has access to unique insights regarding cyberattack trends targeting civil society—a critical pillar of global democracy. In addition, general attack data derived from the company's network—which processes more than 20% of the world's Internet traffic—allows for a comparison between attacks against civil society and other Internet users more broadly.

Cloudflare identified four material attack trends over the past year: DDoS attacks aimed at taking sites offline, attempts to exploit security websites vulnerabilities, phishing attacks, and Internet shutdowns. The data demonstrates how civil society organizations were targeted more frequently, and often more intensely, than other Internet users. These findings reinforce previous Cloudflare research, which found that religious institutions, non-profits, and other civic groups were some of the most targeted by cyberattacks in 2025.

In addition to attacks for financial gain, civil society organizations are uniquely targeted due to the nature of their work. The consequences can be dire. Successful intrusions can expose sensitive data—such as the identity of confidential sources or location of activists—that can enable surveillance, prosecution, or targeted violence. A DDoS attack can upend the provision of life-saving aid during a natural disaster. A phishing campaign can siphon limited donor funds that an organization cannot replace.

Cyberattacks are one of many challenges currently facing civil society organizations. In addition to navigating what experts describe as a general shrinking of civic spaces around the world, non-profit leaders also report difficulties with decreased government funding, an uncertain economic environment, and greater demand for services.

As a result, many organizations report operational and financial strain, reduced staff, and shuttered digital security programs. In 2025, NetHope reported that fewer than one-third of non-profits considered their cybersecurity budgets adequate. Moreover, the rapid development and adoption of artificial intelligence (AI) adds another layer of complexity to the security environment, accelerating existing threats while also offering new tools for defense.

This report is intended to help inform civil society, policymakers, and the public about the security challenges facing these organizations and the importance of promoting broad access to cybersecurity services.

Methodology

This study assesses technical threats against Project Galileo participants during a one-year period from February 1, 2025, to January 31, 2026. Participants include a diverse range of civil society organizations, broadly defined to include independent, non-profit entities that operate to advance the public interest. Groups are further categorized based on the primary nature of their work: media, human rights, environmental, and social welfare.

To protect the privacy and security of participants, Cloudflare analyzed anonymous and aggregated web traffic and email data to identify broad trends and patterns. For application-layer DDoS attacks and website vulnerabilities, Cloudflare reviewed 2,801 of the more than 3,400 Internet properties with active traffic during the report's coverage period, including sampled web traffic. The report also used a representative fraction of total requests to analyze both legitimate and malicious activity. Data regarding phishing attacks was obtained from a smaller subset of more than 70 participants that utilized Cloudflare Email Security. Internet shutdown data is identified from overall web traffic that traverses across the company's global network. In addition to this quantitative data, Cloudflare published case studies exploring the security needs of 16 organizations and how they are addressed. Cloudflare received explicit permission from organizations mentioned in this report.

Mapping cyberattacks against civil society

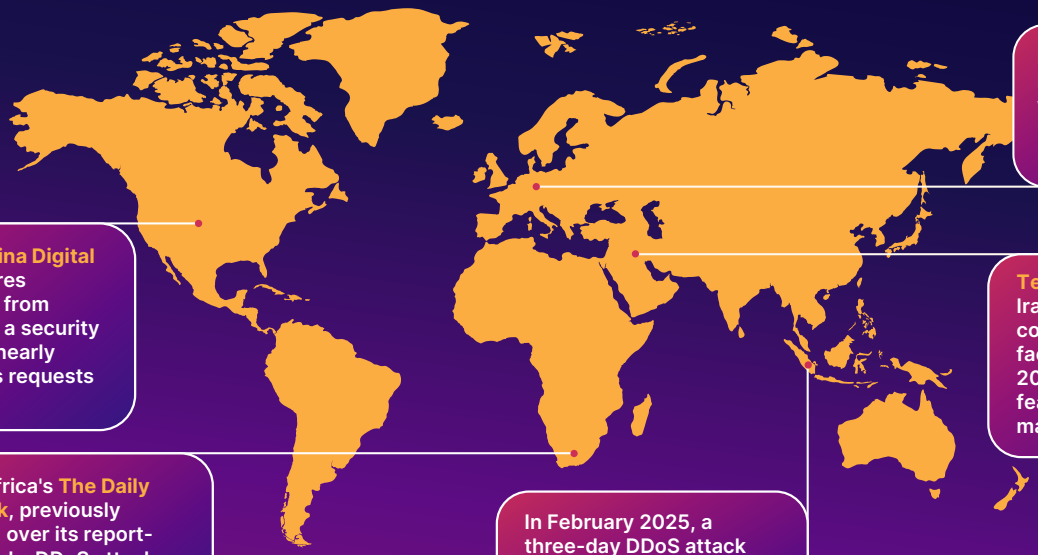

Médecins Sans Frontières in Germany was hit by a high-volume DDoS attack lasting an hour and 41 minutes.

The US-based **China Digital Times**, which shares uncensored news from China, introduced a security rule that blocked nearly 21,000 suspicious requests in one day.

Tech4Peace, an Iraq-based organization countering digital threats, faced five DDoS attacks in 2025, one of which featured 2.6 billion malicious requests.

South Africa's **The Daily Maverick**, previously targeted over its report-ing, faced a DDoS attack peaking at 165,000 rps.

In February 2025, a three-day DDoS attack targeted **Wahana Visi Indonesia**, a Christian humanitarian organization.

Malicious requests against civil society mitigated by Cloudflare

Africa	Asia	Europe	Middle East	North America	Oceania	South America
8b	7b	17b	780m	5b	112m	639m

Distributed denial-of-service attacks

Application-layer DDoS attacks were the most common form of cyberattack impacting civil society organizations. They accounted for 31.43 billion, or 81.7%, of the total 38.5 billion malicious requests recorded during the reporting period. Application-layer attacks overwhelm a site's application, web pages, or APIs, by flooding it with seemingly legitimate requests that drain its resources, slowing or knocking the site offline entirely.

More than 10% of all traffic to human rights organizations was classified as a part of a DDoS attack. This was the largest share of any group of participants, and roughly 40 times more than social welfare organizations.[1] Media groups were widely targeted, with more than one in eight facing malicious DDoS traffic during the reporting period.

What is a DDoS attack?

A distributed-denial-of-service (DDoS) attack is a malicious attempt to disrupt the normal traffic of a targeted server, service, or network by overwhelming the target or its surrounding infrastructure with a flood of Internet traffic. Over time, DDoS attacks have adapted into highly versatile threats capable of targeting the network, transport, or application layers of the Internet with a variety of attack categories and tactics.

Application layer attacks attempt to exhaust backend server resources by directly targeting the applications and services responsible for processing web requests. For example, HTTP flood attacks utilize HTTP GET or POST requests to overwhelm a targeted server or application by flooding a specific URL or network port (the designated communication endpoint through which a server receives web traffic). Because HTTP is the foundational protocol utilized to load webpages, these malicious requests closely mimic legitimate user behavior. Consequently, the target's resources become exhausted resulting in severe performance degradation or failure.

Attackers often combine high-rate attacks like HTTP floods with other tactics resulting in multi-vector campaigns. For example, Slowloris is an application-layer attack that operates by utilizing partial HTTP requests. An attacker overwhelms a targeted server by opening and maintaining numerous HTTP connections and keeping them open for as long as possible. When the server's maximum possible connections have been exceeded, each additional connection will not be answered and denial-of-service will occur. Unlike volumetric attacks that require lots of machines and bandwidth, Slowloris falls into the category of "low and slow" attacks, which require limited bandwidth and can be successfully launched by a single machine.

To learn more about network and transport layer attacks, such as SYN floods, UDP flood, and PING (ICMP) floods, as well as DDoS tactics and tools like Smurf attacks and High Orbit Ion Cannons (HOIC), visit the Cloudflare Learning Center.

1. Cloudflare generally categorizes Project Galileo participants into the following categories: human rights defenders, social welfare organizations, journalism organizations, and environmental groups.

DDoS attacks against civil society were longer in duration than other attacks Cloudflare mitigated during the time period. In 2025, approximately three quarters of all the application-layer attacks Cloudflare saw ended within 10 minutes. For civil society specifically, the opposite was true. Nearly every one of the largest attacks lasted longer than 10 minutes, with some spanning days and weeks. Additionally, by the company's standards, each attack against civil society was small-to-medium in size, with peaks ranging from roughly 13,000 to 813,000 requests per second (rps). These sizes were consistent with the application-layer DDoS attacks Cloudflare mitigated for all its customers.

The Christian humanitarian organization Wahana Visi Indonesia faced a three-day long attack in February 2025. Spread over 13 bursts of activity, the attack featured 4.9 billion malicious requests, with a peak of 366,666 rps, making it one of the largest against civil society during the coverage period. Similarly in April 2025, the Refugee Council, a British organization supporting refugee integration, faced a DDoS attack that spanned seven days and 15 hours. Compared with the Wahana Visi Indonesia attack, this one was smaller in scale with a total of 261.3 million malicious requests spread over 14 chunks of activity, peaking at 18,333 rps.

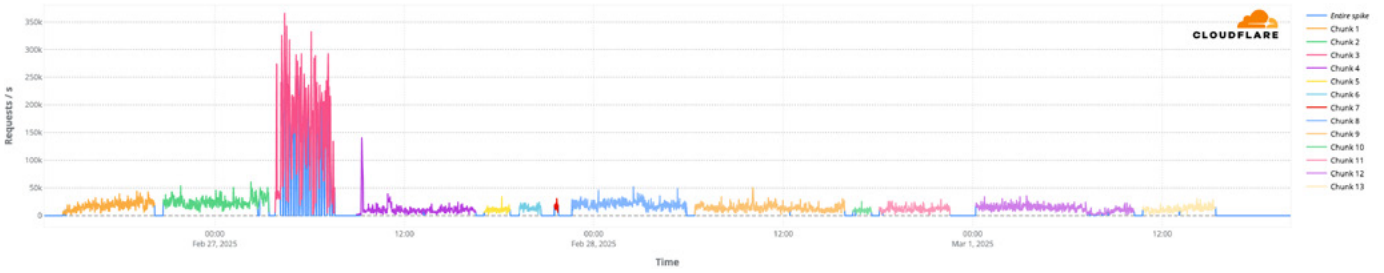

DDoS attack against an wahanavisi.org in February and March 2025

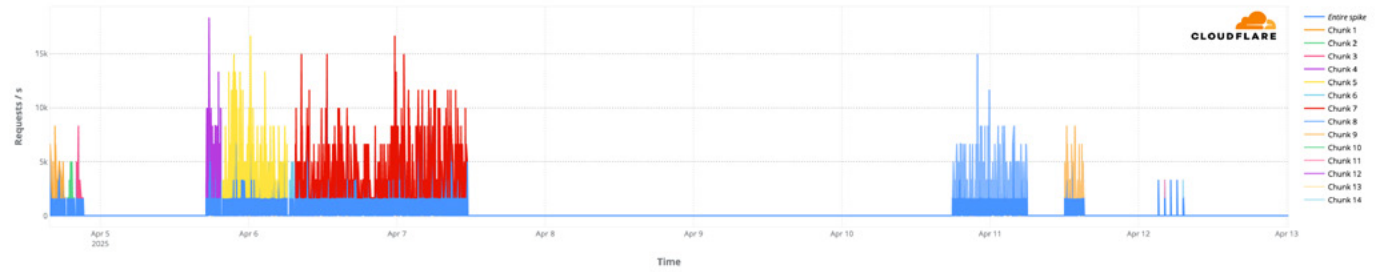

DDoS attack against the refugeecouncil.org.uk in April 2025

The extended duration and chunked structure of these attacks signal a deliberate effort to wear down the websites' defenses over time. By sending malicious traffic in short bursts separated by intentional pauses, attackers were able to fall out of scope of mitigation defenses before resuming the attack. The pauses let attackers reverse engineer defensive rate limits, see which rules were triggered, and then adjust their traffic signatures. They also allow attackers to exploit dynamic fingerprint rules, short-lived defenses built to match an attack's specific patterns that expire once traffic stops so legitimate visitors are not blocked. Pausing long enough for those rules to fade forces the mitigation system to detect each new wave of activity from scratch.

In contrast to these longer campaigns, a Dutch LGBTQ+ rights organization faced a 15-minute attack that featured only one burst of malicious activity. The volume of malicious traffic was still high, peaking at 375,000 rps and reaching a total of 99.9 million requests. The attack occurred in May 2025, three days after the International Day Against Homophobia, Transphobia and Biphobia.

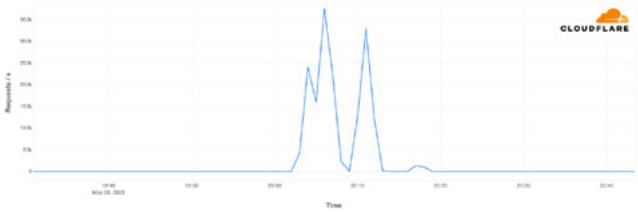

DDoS attack against a Dutch LGBTQ+ rights organization in May 2025

How to protect against a DDoS attack

DDoS mitigation services sit between a website and incoming traffic as a reverse proxy, analyzing requests before they reach the origin server. Providers like Cloudflare operate a global network spanning hundreds of cities, distributing attack traffic across that infrastructure so no single server is overwhelmed. These mitigation tools automatically detect and filter out malicious requests regardless of what content the site is hosting, while letting legitimate visitors to the site.

Cyberattacks against independent media in exile

Media groups in exile rely on the Internet to reach communities in their origin countries. But they face attacks aimed at preventing them from communicating with those audiences. Nearly 5% of the 41 billion requests to journalism-in-exile sites were malicious, almost four times the rate seen across journalism organizations overall.

elTOQUE

Run by journalists in exile, elTOQUE provides Cuban audiences with high-quality, independent journalism. In December 2025, the outlet's website faced a DDoS attack featuring nearly 426.8 million malicious requests and peaking at 108,167 rps—a medium-sized attack by Cloudflare standards. The outlet believes the attack may be linked to its tool for comparing exchange rates between the Cuban peso and foreign currencies, which the Cuban government has referred to as "economic terrorism." During the same month as the attack, elTOQUE's website was blocked in Cuba.

The Moscow Times

In July 2025, The Moscow Times' website experienced a DDoS attack featuring 123.4 million malicious requests, peaking at 319,000 rps. The outlet, which relocated to Amsterdam after the Russian military's invasion of Ukraine, was designated as "undesirable" in 2024, effectively banning its operations in Russia and putting staff at risk of criminal prosecution. Because the organization remains in exile, operating securely online is fundamental to its ability to provide independent reporting to the Russian public.

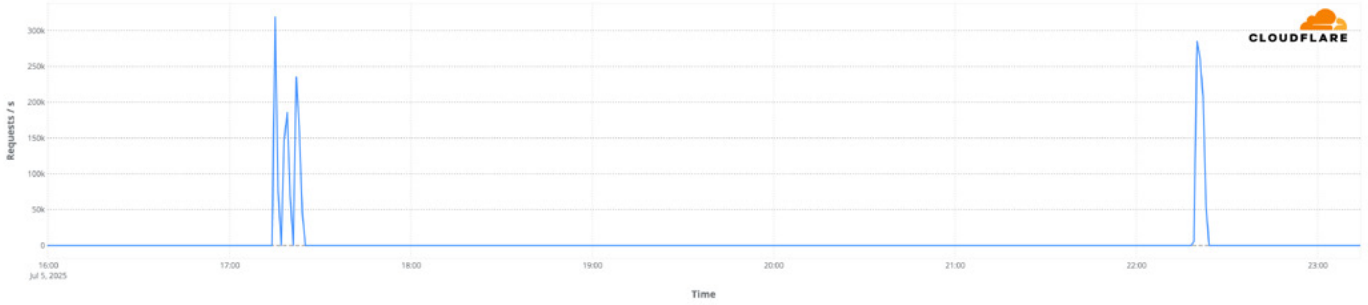

DDoS attack against themoscowtimes.com in July 2025

> **As a resistance media dedicated to preserving the memory of the Chinese Internet, China Digital Times confronts relentless attacks every single day."**
>
> **Xiao Qiang**
> China Digital Times

Website vulnerabilities

Civil society organizations faced attempts to exploit website vulnerabilities at a rate more than seven times higher than other Cloudflare customers. Website vulnerabilities are a type of cyber attack that targets flaws in outdated or unpatched systems, that allow threat actors to extract sensitive data or access internal systems.

According to Cloudflare data, journalists were disproportionately targeted with this type of attack. Media organizations faced 40.5% of the 7.1 billion attempts Cloudflare mitigated despite only accounting for 22.7% of participants. That equates to an average rate of 4.5 million maclicious requests per organization, or roughly one every seven seconds.

Cloudflare identified three dominant ways that threat actors probed websites for weaknesses. The most common, accounting for 44% of probing activity, was HTTP anomalies, which are malformed requests sent in bulk to trick security systems like a Web Application Firewall into allowing malicious payloads to pass through. For example, in September and October 2025, the managed rules on Cloudflare blocked 162.3 million requests against a US-based publishing non-profit. 99.9% of the requests were HTTP anomalies.

The next two most common techniques were SQL injection (16%), attempts to manipulate a site's database into handing over private data, and vulnerability scanners (15%), the use of automated tools looking for outdated software with known weaknesses. All three methods were overwhelmingly conducted by bots, allowing attackers to conduct scanning at a scale and speed humans cannot match.

Average probes for security weaknesses, per organization

Media organizations saw the highest rate of probing per organization

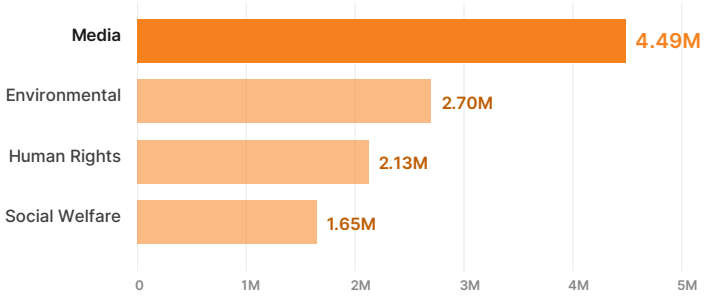

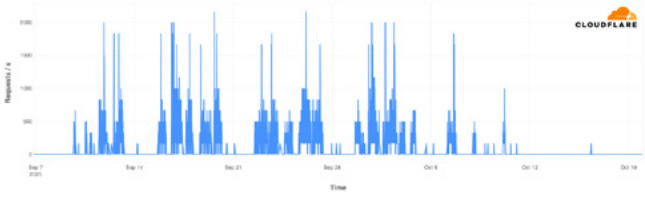

Web vulnerability attack against a US non-profit publishing organization in September and October 2025

What is a web application vulnerability?

A web application vulnerability is a weakness in a website's code, configuration, or underlying software that an attacker can exploit to steal data or disrupt operation. They do so by sending specially crafted requests, often disguised as ordinary user input like a login form submission. The site then behaves in ways its developers never intended, such as exposing private data, granting unauthorized access to user accounts, or letting attackers take control of the site.

Web application vulnerabilities range from common bugs in a site's code to sophisticated, targeted exploits designed to bypass traditional defenses. Attackers continuously scan the Internet for these weaknesses, often using automated tools that test thousands of sites per minute.

To learn more about common web application vulnerabilities such as SQL injections, cross site scripting and Zero-day vulnerabilities, visit the Cloudflare Learning Center.

For journalists and civil society organizations, unauthorized access to internal systems can cause not only operational disruption, but also potentially jeopardize the safety and security of employees, sources, and other vulnerable populations. For example, China Digital Times, which shares uncensored news and online voices from China, reported that exposure of their editors' identities would put them at risk of harassment, arrest, and other offline repercussions when traveling to China. The organization used specialized security tools to protect their website from malicious bot traffic probing for how to infiltrate their systems. For example, one rule blocked more than 99.9%—all but eight of 21,000—requests in 24-hours in November 2025.

Cloudflare data also indicated that web vulnerability attacks were carried out alongside other types of attacks. For example, threat actors appeared to use high-volume DDoS attacks to mask simultaneous scanning for website vulnerabilities. In November and December 2025, a global environmental organization experienced this type of multilayered attack in the same month as a major climate conference in Brazil. Specifically, Cloudflare mitigated 3.34 million malicious vulnerability scanning attempts, as well as a DDoS attack featuring 6.97 billion requests with a peak of 118,333 rps.

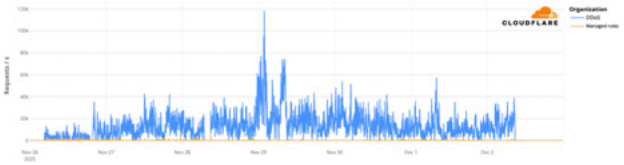

Simultaneous web vulnerability and DDoS attacks against a global environmental organization in November and December 2025

How to protect against website vulnerabilities

Protecting against a website vulnerability requires a layered approach. Key steps include enforcing multi-factor authentication, using an SSL certificate to encrypt traffic, keeping software and plugins up to date, and running proactive security scans.

Another critical layer of defense is a web application firewall (WAF), which sits between a website and incoming visitors to inspect every request before it reaches the server. WAFs compare incoming requests against rules designed to identify malicious patterns, such as a database query hidden inside a search box or a script injected into a form field, and block anything that looks dangerous before it can reach the application. Modern WAFs also use machine learning to identify unusual patterns of behavior, such as a single visitor making thousands of requests in a few seconds or probing the site for known weaknesses, and adapt their rules in real time as new attack techniques emerge. Many providers offer these managed rulesets, including Cloudflare, that automatically learns and adapts to keep web applications safe based on new threat intelligence.

Phishing attacks

Nearly 10 percent of the approximately 29 million emails Cloudflare processed for civil society included potential phishing material. Compared to other customers, civil society organizations were targeted with a higher concentration of emails that contained harmful links, attachments, and other content intended to gain unauthorized access or steal sensitive information.

Cloudflare identified three common phishing methods deployed against civil society, which aligned with patterns the company observed more broadly, including in an analysis of 450 million emails. First, 19.5% of emails identified as potential attacks used deceptive URLs to lure users into visiting malicious websites. Second, 16.8% of these messages impersonated a specific individual that the receiver recognized and viewed as a trusted authority. Finally, 13.4% impersonated a trusted brand name, such as through spoofed domains or unauthorized logo usage. The top five most frequently impersonated brands were, in descending order, Apple, Docusign, Datadog, American Express, and Intuit. In addition, nearly half of the phishing threat emails Cloudflare identified used newly registered domains. Without an established history, these domains can evade security features that identify known phishing sites.

Cloudflare data also confirmed that phishing attacks continue to increase in sophistication. For example, Cloudflare Email Security identified 1.2 million highly malicious emails targeting civil society. Nearly a third (30.2%) of those emails were able to bypass standard authentication checks that rely on signatures like sender, origin, and content integrity. Cloudflare Email Security did ultimately flag these emails as malicious, however, the share that would not have been identified by basic email security measures reflects the growing sophistication of these campaigns.

What is a phishing attack?

A phishing attack is the use of deceptive messages from seemingly reputable sources to trick victims into revealing sensitive information like login credentials, passwords, or financial data for malicious use.

Attackers typically utilize email, text messages, or malicious websites featuring deceptive URLs and unauthorized logos to trick users into clicking links or opening attachments. Such actions can lead to fake login pages that steal credentials or install malware to capture passwords and session cookies.

To learn more on common phishing attacks, visit the Cloudflare Learning Center.

Threat categories within civil society emails
Cloudflare identified 2.8 million emails that potentially contained phishing material.

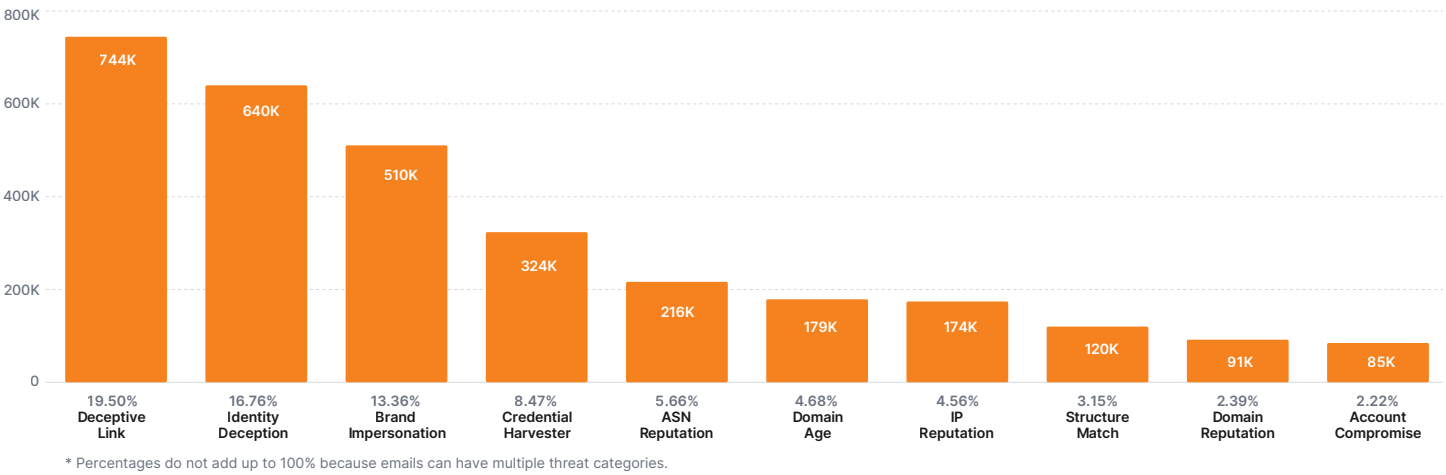

* Percentages do not add up to 100% because emails can have multiple threat categories.

Phishing attempts have also grown more personalized. Citizen Lab reported in April 2025 that a phishing campaign targeting the World Uyghur Congress (WUC) was unique not for its technical sophistication but for the attackers' deep understanding of the exiled Uyghur community. Threat actors used impersonated emails to trick WUC into downloading a trojanized version of a trusted Uyghur-language text editor, seeking to install Windows malware to remotely surveille the organization. This attack relied on WUC trusting the text editor brand to gain access. These types of phishing attacks are more difficult to mitigate using automated detection tools because they rely more on social engineering tactics rather than technical sophistication.

AI may exacerbate the effectiveness of these campaigns, as large language models allow threat actors to generate hyper-customized and realistic content at unprecedented speed and scale. A March 2026 Huntress investigation highlighted this shift, identifying an expansive campaign that was suspected of using AI-generated emails to target Microsoft cloud accounts across more than 340 entities, including civil society organizations.

How to protect against a phishing attack

Phishing can occur across various attack vectors, but email remains the most common. Email security tools scan incoming messages to screen things like senders, links, and attachments. To verify senders, these tools check digital signatures against a set of records published by legitimate organizations; any mismatch or unauthorized origin flags the message. To evaluate links, URLs are cross-referenced with databases of known malicious sites and sometimes tested in isolated environments. Attachments undergo similar scanning for malware within a sandbox. Finally, tools analyze message content for phishing markers like false urgency, employee impersonation, or unusual requests, subsequently blocking or quarantining identified threats. Learn more about the common signs of a phishing attack.

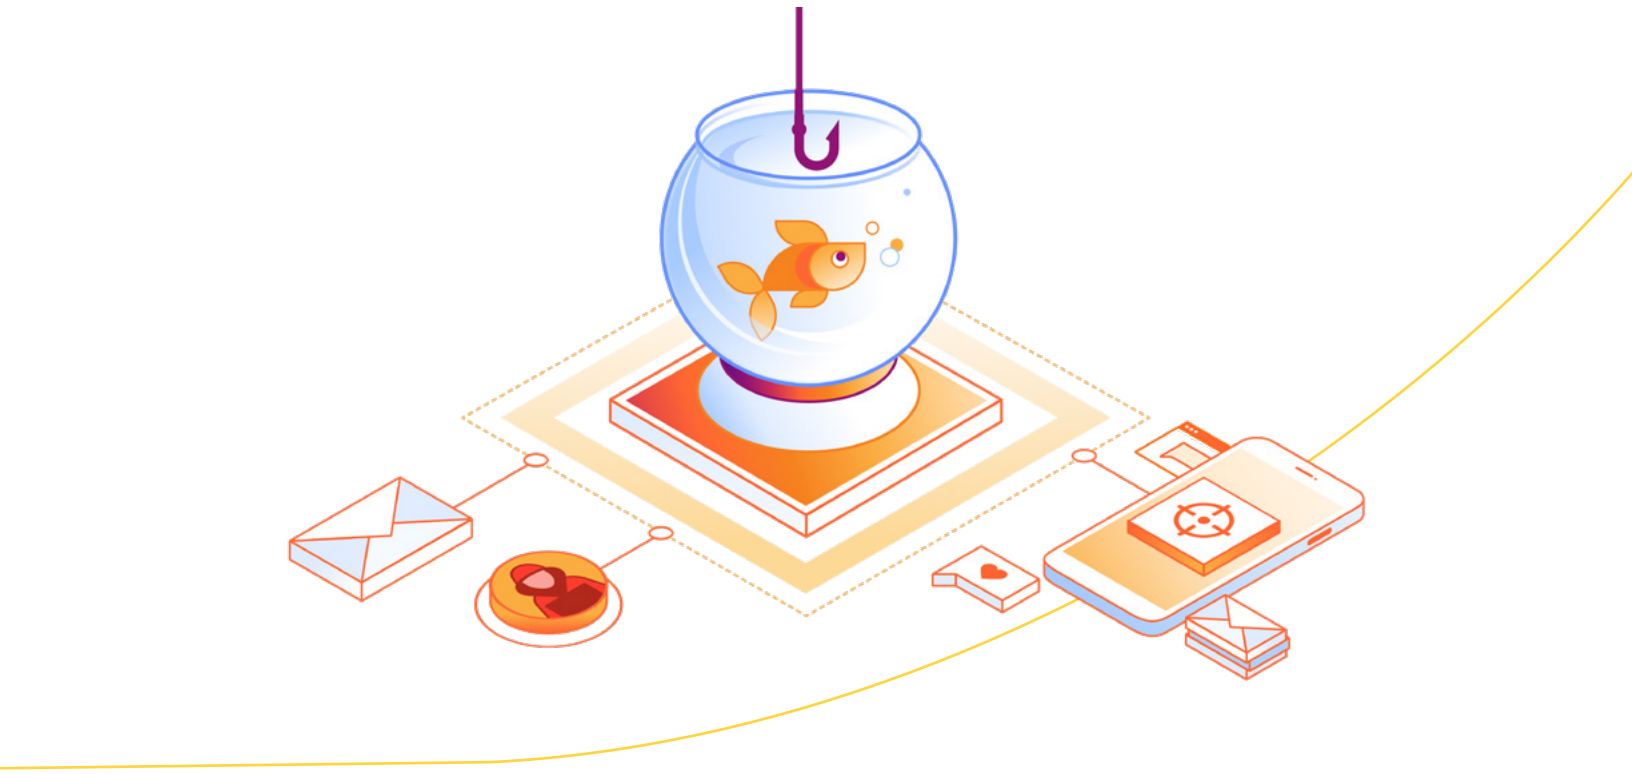

Cyberattacks against organizations supporting digital rights

Organizations working to advance a free and open Internet are frequently targeted for disruption. More than a third of all traffic to digital rights organizations was malicious. Cloudflare mitigated 7.1 billion such requests, an average of 19.5 million per day.

Tech4Peace, an Iraq-based organization countering digital threats across the Middle East, was hit by five DDoS attacks across four months in 2025. Together, these attacks accounted for the majority of the malicious traffic impacting digital rights organizations. Multiple attacks against Tech4Peace aligned with the timing of high-visibility publications. For example, in May 2025, the organization was hit by an eight-day DDoS attack after it published an article debunking an AI-generated image of a Syrian politician bowing to US President Trump. The attack featured more than 2.6 billion malicious requests across 15 distinct chunks, highlighting a volatile and sustained effort to disrupt the organization's online operations. A month before, a separate attack occurred within a day of an article that drew more than 156,000 views, far above the organization's typical readership.

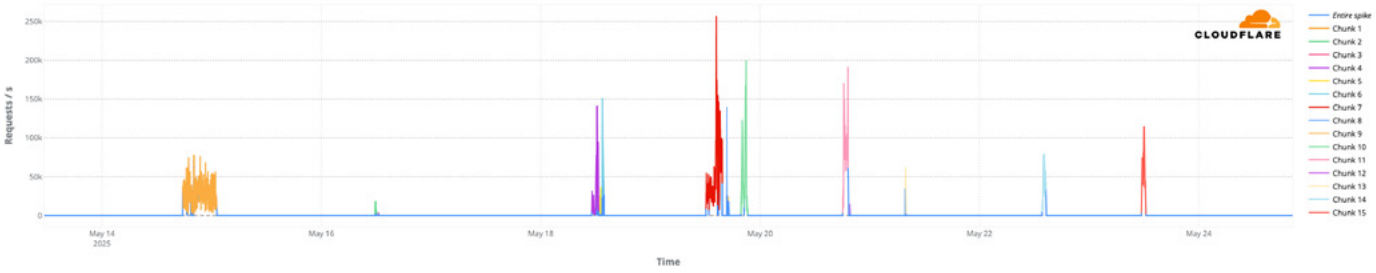

Eight-day DDoS attack against t4p.co in May 2025

> **Cybersecurity tools protect our infrastructure and enable us to focus on our everyday activities. This often involves working with journalists and activists who are targeted by repression. We have to protect both our staff and the beneficiaries of our work."**
>
> **Bojan Perkov**
> SHARE Foundation

> **Human rights defenders and journalists face a disproportionate share of online threats. In the Asia-Pacific region, where many of the organisations we support operate in constrained or hostile digital environments, that risk is acutely felt."**
>
> **Khairil Zhafr**
> EngageMedia

Internet shutdowns

Cloudflare reports on Internet shutdowns detected across the company's network to assist human rights defenders and others in civil society documenting these disruptions. During the reporting period, 183 Internet disruptions were identified, 85 of which appeared to be government-directed based on public reporting. These intentional shutdowns occurred during student exam periods, protests, elections, and armed conflict. Such restrictions limit the ability of civil society and media to disseminate independent information, provide public services, communicate with communities at home and abroad, and mobilize for change. Organizations like the Internet Society have also sought to estimate the significant economic cost of shutdowns.

Prior to its January 15 general election, the Uganda Communications Commission ordered Internet service providers to restrict public Internet access and certain mobile services. Cloudflare data identified a subsequent 95% drop in traffic within 30 minutes of the order going into effect. Justified under an effort to curb "misinformation, disinformation, electoral fraud and related risks, and preventing incitement to violence," the order followed the suspension of permits for several non-profits. Independent experts have highlighted how the shutdown and permit suspensions caused a chilling effect among civil society, constraining their ability to provide oversight of the electoral process, share independent reporting, and fundraise. The Uganda-based organization CIPESA estimated that the shutdown cost the country's economy $16 million.

Causes of Internet outages globally
From February 2025-January 2026, 15 causes observed

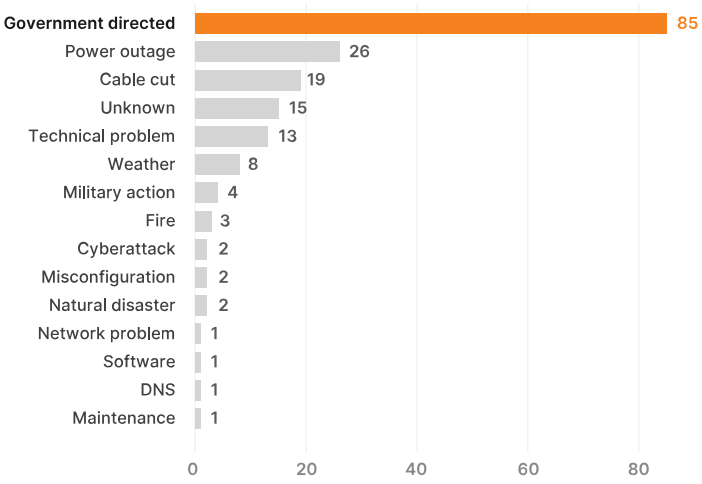

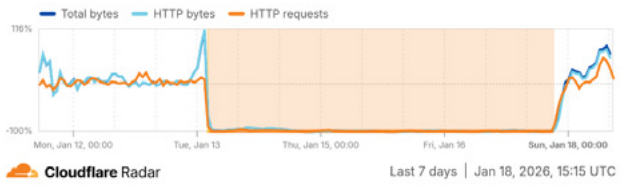

Traffic volume in Uganda during the January 2026 shutdown

> **Internet shutdowns are a direct attack on civic space. For civil society organisations, they disrupt communication, block documentation and monitoring, weaken emergency response, and isolate the communities we serve. Uganda's election shutdown, like similar restrictions elsewhere, demonstrated that cutting connectivity is not merely a technical measure, but a serious interference with rights, accountability, and democratic participation."**
>
> **Patricia Ainembabazi**
> CIPESA

In Iran, Cloudflare identified eight government-directed Internet shutdowns during this report's coverage period. Starting around 16:30 UTC on January 8 and within 30 minutes, traffic in the country fell rapidly by nearly 90%. It dropped to effectively zero by 18:45 UTC, signaling a complete shutdown in the country and disconnection from the global Internet. Civil society organizations reported that the restrictions, which lasted until February 1, limited their ability to document the government's use of force against protestors.

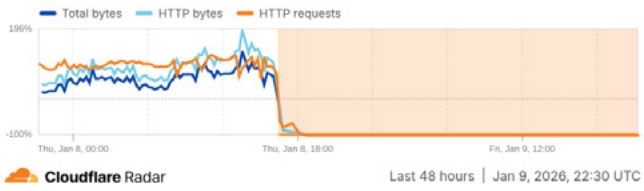

Traffic volumes in Iran during the January 2026 shutdown

Supporting the Internet measurement community

Cloudflare shares its data on Internet outages directly with civil society, supporting their independent work to monitor and combat shutdowns. Projects that incorporate Cloudflare data include the Internet Society's Pulse project, Access Now's #KeepItOn campaign, and the Open Observatory of Network Interference's Explorer.

Conclusion

Cloudflare research consistently shows that cyberattacks are increasing across the Internet in both frequency and sophistication. This report demonstrates that across numerous threat types—DDoS attacks, website vulnerabilities, and email phishing—civil society organizations were targeted more frequently, and often more intensely, than other Internet users. These findings align with similar research from organizations such as NetHope and Protect.ngo, which together underscore the scope of the security challenges facing non-profits.

Drawing on both quantitative data and qualitative case studies, this report also reveals a common pattern. Cyberattacks against civil society organizations often coincided with critical moments in their work, such as publishing investigative reporting, navigating elections, or conducting public advocacy. This reinforces a well-documented reality: civil society organizations are targeted not only because of the work they do, but also to disrupt the moments when that work is most impactful.

This report was produced alongside the rapid development and adoption of artificial intelligence. AI is expected to accelerate existing threats and introduce new types of attacks. Because civil society organizations remain disproportionately targeted by cyberattacks, they are likely to experience additional security consequences as a result of their increased AI adoption. However, AI will also provide significant opportunities to improve security by detecting behavioral anomalies, identifying data leaks, and automating and scaling mitigation. If adopted responsibly, AI can materially improve civil society's cybersecurity.

Recommendations

Based on the research, findings, and case studies prepared for this report, Cloudflare makes the following recommendations.

Promote universal access to cybersecurity services. Given the ubiquity of attacks, cybersecurity is essential to meaningful participation online. These services should be made simple to use and affordable for everyone, including civil society. Relatedly, efforts to limit the availability of cybersecurity services should be weighed against their impact on all Internet users. Each organization that strengthens its defenses contributes to a more secure network, while each one left exposed becomes a conduit for attacks against others.

Increase transparency about cyberattacks and Internet shutdowns. Everyone should be able to understand how the Internet works. Expanded research, monitoring, and analysis of cyberattacks and Internet shutdowns can shed light on the actors behind them and how tactics evolve over time. The resulting data can be shared among a broad set of stakeholders and, where possible, made public. Such research and stakeholder coordination can inform effective policy, legal, and technical solutions to emerging threats.

Improve the accessibility of AI-enabled security services. As AI becomes increasingly adopted by civil society organizations and leveraged in cyberattacks, access to the most advanced security tools—such as AI-enabled defenses and post-quantum cryptography—is critical to avoiding a widening technical gap. The strongest protections should be integrated into tools by default, and barriers to obtaining them should be lowered. The result will be a higher baseline of security for everyone, including civil society organizations, against evolving threats.

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

REV:BDES-9187.2026JUN16