CLOUDFLARE®

Five Best Practices for Mitigating DDoS Attacks

How to defend against rapidly evolving Distributed Denial-of-Service threats and address vulnerabilities at every layer

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

CLOUDFLARE

I. Executive Summary

Distributed denial-of-service (DDoS) attacks remain one of the most effective methods used by cybercriminals to cause significant financial, operational, and reputational damage to businesses worldwide. Though these attacks take different forms, the goal is always to incapacitate targeted servers, services, or networks by flooding them with traffic from compromised devices or networks.

As organizations have hardened their defenses, cybercriminals have responded with newer attack types targeting multiple applications and services. Some of these attacks target layers 3 and 4 of the Open Systems Interconnection (OSI) model in new ways, resulting in network traffic spikes of up to 1.3 TB per second or more. Others are low-speed, low-intensity Layer 7-based offensives designed to fly under the radar and target one or more services gateways and application layers.

Meeting the challenges associated with DDoS attacks requires a comprehensive approach which addresses all threats at all layers. But enhanced security shouldn't come at the expense of performance. While on-premise solutions can be part of the answer, a more robust solution will integrate performance with scalable, cloud-based mitigation that works at the network edge to deliver maximum agility and unlimited capacity.

PART 1

What is a DDoS attack?

A distributed denial-of-service (DDoS) attack is a malicious attempt to disrupt normal traffic flowing to a targeted server, service, or network by overwhelming it with a flood of Internet traffic. To be effective, these attacks require threat actors to take control of online computers, routers, IoT devices or other endpoints to leverage as sources of attack traffic. These machines are infected with malware and then weaponized in a "botnet" that is activated by remote control.

When the IP address of a targeted server or network is established, each bot sends simultaneous requests to that target with the intention of pushing it to overflow capacity, resulting in a denial-of-service to normal traffic. Since each bot is a legitimate device, separating attack traffic from legitimate traffic can be extremely difficult.

Varieties of DDoS attacks

DDoS attacks can target any of the 7 distinct "layers" within the OSI model for network connections. While all of these attacks involve suffocating targets with malicious traffic, they can be divided into three distinct categories.

VOLUMETRIC ATTACKS: These attacks are designed to create congestion between the site and the larger Internet by targeting the network, slowing down web performance and degrading access for legitimate users.These attacks often employ DNS amplification and other techniques to create massive traffic surges measured in bits per second (Bps).

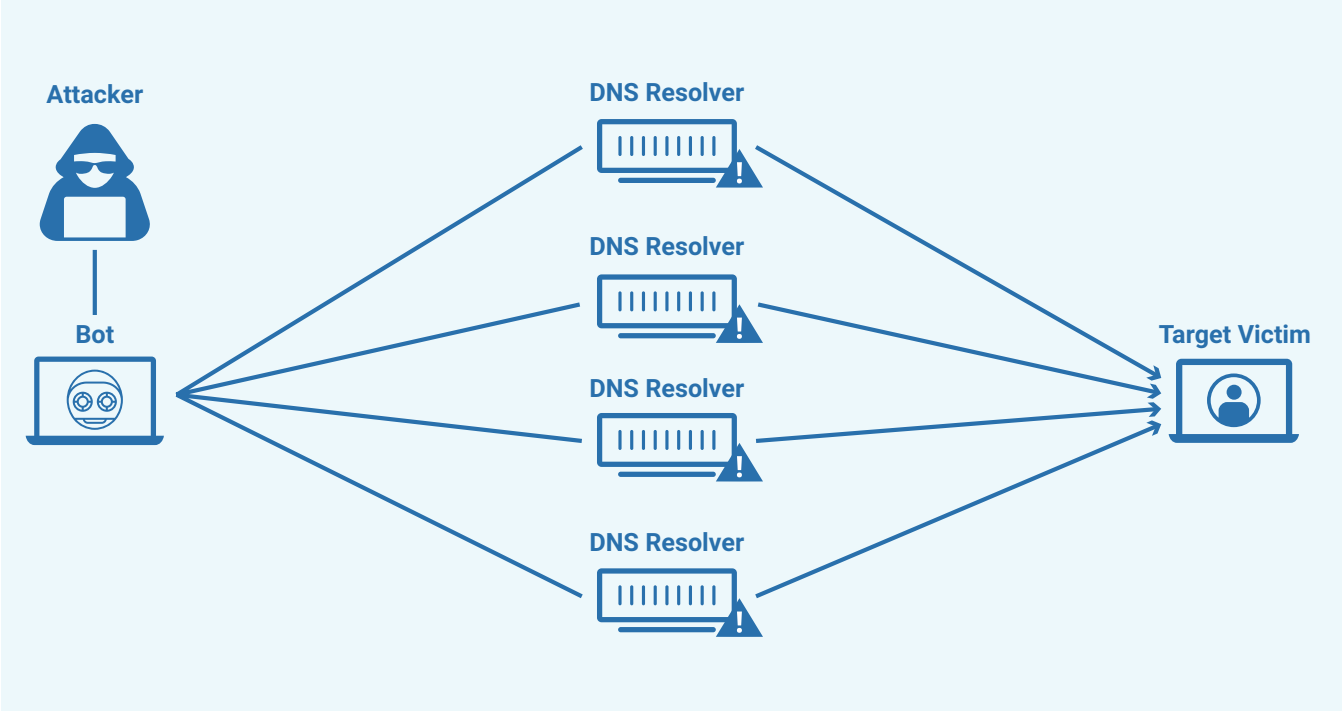

PROTOCOL ATTACKS: The objective of protocol attacks is to target vulnerabilities in Layers 3 (network) and 4 (transport) of the OSI model and consume all the available capacity of web servers or their intermediate resources—including firewalls and load balancers. These attacks can involve SYN floods, Ping of Death attacks, Smurf DDoS, and fragmented packet attacks, which are all measured in packets per second (Pps).

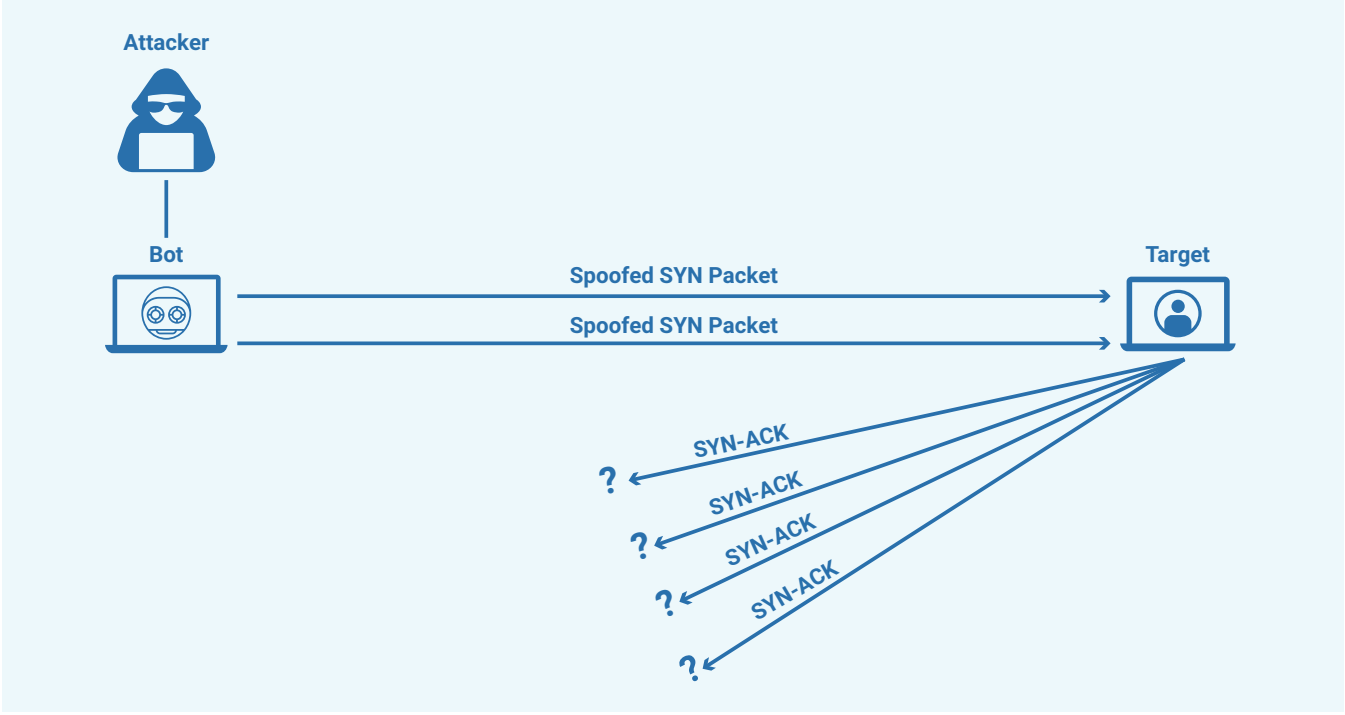

APPLICATION LAYER ATTACKS: Sometimes called Layer 7 DDoS attacks, these attacks target the layer where web pages are generated on the server and delivered in response to HTTP or HTTPS requests. Akin to repeatedly hitting refresh in a web browser on many different computers all at once, the resulting flood of HTTP/S is measured in requests per second (Rps).

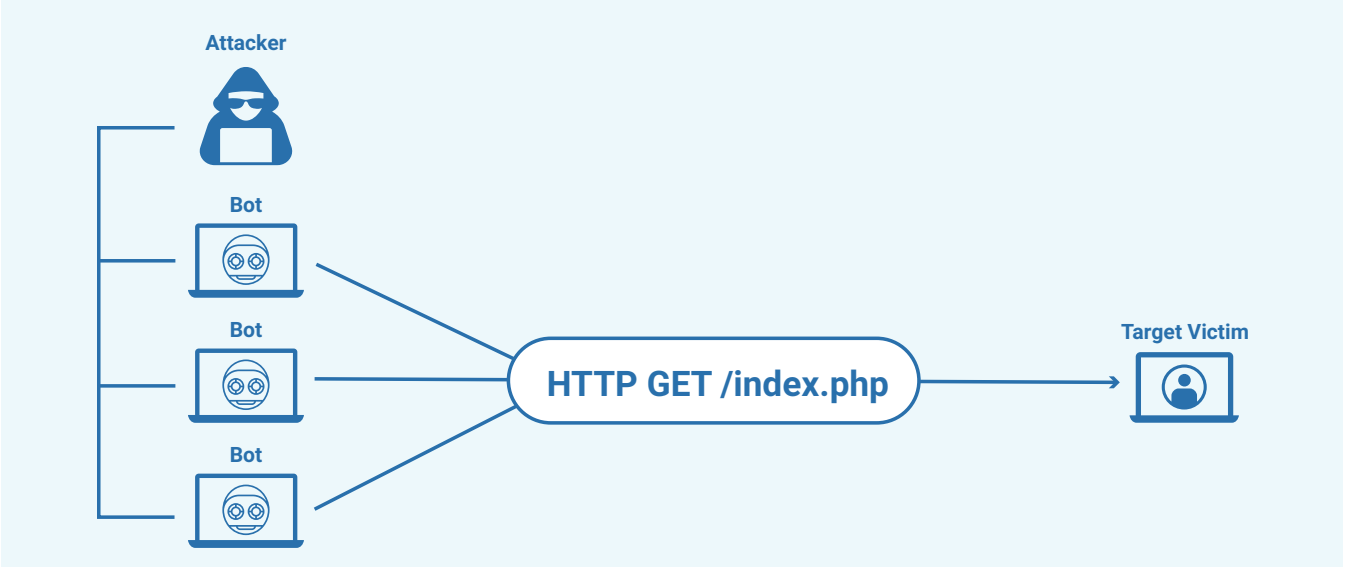

There is some overlap between these types of attacks. Some protocol attacks can be volumetric, for instance. And then there are multi-vector attacks, in which threat actors target multiple layers of the protocol stack at the same time, or cycle attack vectors based on the countermeasures taken by the target. Furthermore, many multi-vector attacks are merely smokescreens designed to provide cover for an attempted data breach or other crime.

How DDoS attacks damage your business

Getting knocked offline by DDoS attacks can have a detrimental impact on revenue, customer service, and basic business functions. Whether the aim is to cripple your site or network, to divert traffic to rivals, top mask the theft of corporate data, or simply to cause maximum reputational damage, users will often put the blame squarely on your business. The average DDoS attack can cost $123,000 for a small company and more than $2 million for a large enterprise[1]. And they just keep coming. The total number of DDoS attacks worldwide is expected to rise from 11.9 million in 2020 to more than 14.5 million by 2022[2].

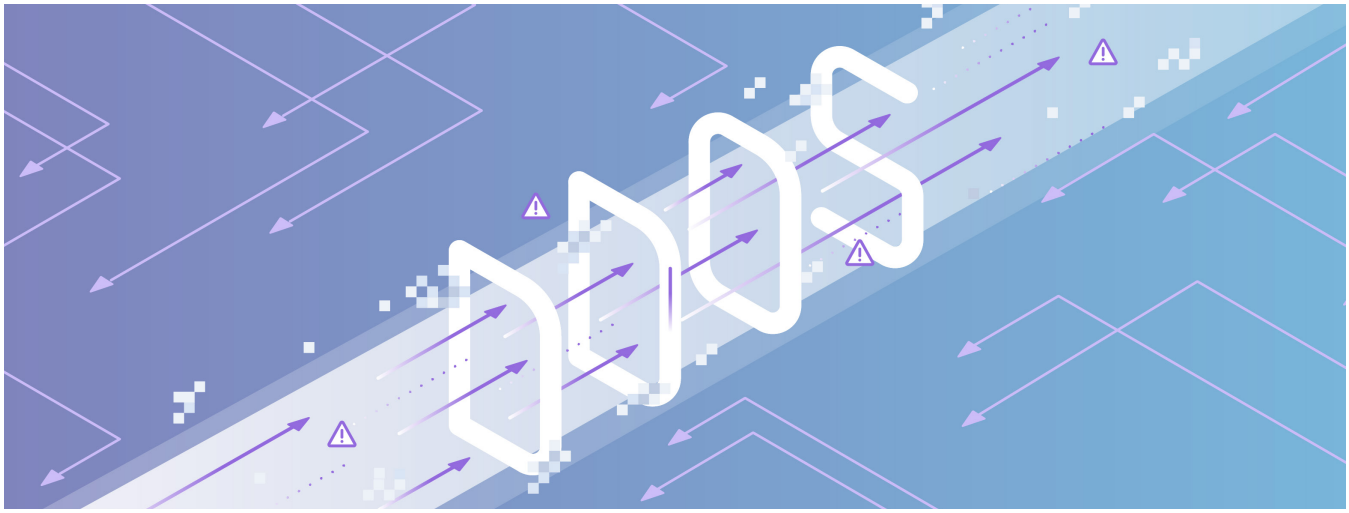

PART 2

Evolving DDoS Challenges and how to address them

Generally speaking, protecting against DDoS attacks requires the ability to:

- Differentiate between traffic spikes stemming from an attack or from high user demand
- Block traffic flowing from botnets without interrupting legitimate traffic
- Intelligently route remaining traffic by breaking it into manageable chunks to prevent denial of service
- Continuously analyze traffic for malicious patterns that can aid in developing adaptive, hardened defenses

Today, two emerging trends are making all of this more challenging.

Volumetric attacks are growing larger

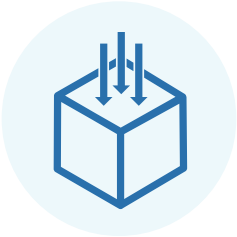

- Volumetric DDoS attacks larger than 100Gbps skyrocketed 967% between 2018 and 2019[3]
- DDoS attacks as large as 1.3 TB per second, like the attack that knocked GitHub offline in 2018, have also become commonplace
- In early 2020, one volumetric network layer DDoS attack is reported to have reached 92 Gbps and 10.38 million packets per second (Mpps)[4]
- While most volumetric DDoS attacks last just minutes, some can last hours, and up to 73% of organizations hit by volumetric attacks are targeted again within 24 hours[5]

Attacks are increasing in complexity

- Three-quarters[6] of all DDoS attack target more than one vector
- DNS amplification targeting Layers 3 and 4, coupled with an HTTP/S flood targeting Layer 7, is an example of a multi-vector DDoS attack
- The more complex the attack, the harder it is to mitigate—the goal of the attacker is to blend in as much as possible, making it all the more difficult to differentiate between legitimate and malicious traffic
- Mitigation attempts to drop or limit traffic are of no use if the attack adapts to circumvent this countermeasure

Based on these requirements and evolving trends, here are five best practices for DDoS mitigation organizations to prioritize.

PART 3

Best practices for DDoS mitigation

1 TAILOR TACTICS FOR WHAT YOU'RE PROTECTING

If your objective is to protect web servers, a reverse proxy will prevent attackers from being able to identify and target your servers' IP addresses. Instead, they will only be able to target the reverse proxy. For more complex Layer 7 DDoS attacks, a web application firewall (WAF) can act as a reverse proxy to shield targeted servers from certain types of malicious traffic. Some companies build or deploy their own reverse proxies, but this requires intensive software and engineering resources, as well as a significant investment in physical hardware.

One of the easiest and most cost-effective ways to realize the benefits of a reverse proxy is to use a content delivery network (CDN). Look for a CDN with global server load balancing, so that your site can be distributed on several servers around the globe. That way, DDoS attacks will be mitigated closer to the source without impacting performance.

If the goal is to protect network infrastructure, Border Gateway Protocol (BGP) rerouting can be used to redirect traffic to scrubbing centers where malicious traffic can be filtered out. That said, rerouting all traffic to a limited number of geographically distant scrubbing centers can add considerable latency.

For this reason, cloud-based DDoS mitigation solutions of sufficient scale are recommended. With cloud-based mitigation, autonomous system numbers (ASNs) are advertised by the mitigation provider, so traffic is routed direct-to-scrub instead of going to the origin server. Traffic is filtered closer to the source of the attack, further reducing latency.

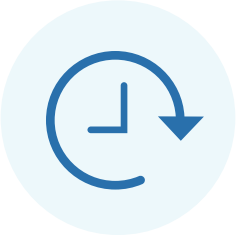

2 PRIORITIZE THE TWO MOST IMPORTANT METRICS—CAPACITY AND TIME-TO-MITIGATION

Assess your existing capacity for mitigating DDoS attacks without impacting site functionality. The traditional approach to absorbing the spikes in traffic generated by DDoS attacks has been to build out on-prem server farms. But this quickly grows costly, and even the most robust enterprise-grade infrastructure is likely to be overwhelmed in the face of volumetric attacks that grow larger by the day.

Rate limiting can help, but it slows down performance and can still result in an outage if your infrastructure is overloaded. When even a few moments of reduced availability can lead to significant lost revenue and productivity, time-to-mitigation (TTM) becomes paramount. To reduce TTM, you'll need to ensure traffic can fail over to an alternate site in the event of an outage—but that will only work for so long before your infrastructure is overwhelmed.

Here again, a more effective approach is to deploy a cloud-based mitigation solution that offers unlimited capacity to protect against DDoS attacks of any scale or complexity, and can provision services at the network edge for maximum agility in mitigating rapidly-evolving DDoS attacks.

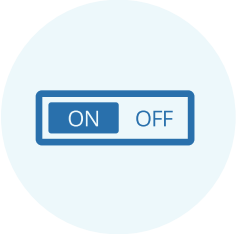

3 CONSIDER ALWAYS-ON VS. ON-DEMAND PROTECTION

With on-demand mitigation services, traffic flows as it normally does until a potential DDoS attack is detected. At that point, traffic is re-routed to the cloud mitigation service, filtered, and passed back to the server of origin. You only pay for DDoS mitigation when it's needed, and no management or additional resources are required. But there are tradeoffs, specifically around time-to-mitigation. Stopping the attack takes longer because traffic spikes must reach certain thresholds before analysis begins and someone manually turns on the mitigation service.

By comparison, always-on mitigation continuously routes and filters all site traffic, so only clean traffic reaches the customer's servers at all times. While more expensive than on-demand services, always-on mitigation provides uninterrupted protection, and leads to faster response times since the service never needs to be turned on manually. What's more, given the growing number of DDoS attacks, always-on services with flat-rate pricing may actually prove less expensive for organizations facing a constant barrage of attacks.

4 NEVER SACRIFICE PERFORMANCE FOR SECURITY

DDoS attacks cause sluggishness and outages that not only degrade performance, but also damage an organization's ability to achieve sustainable growth. Today's digital consumer expects websites and applications to load instantaneously and to never (ever) be offline. Latency becomes noticeable to the average user at 30 milliseconds, and just one second of additional load time can cause conversions to drop by 7%[7].

Latency also damages productivity. Even in the best cases, the average employee wastes one week annually waiting for their network to respond[8]. For Fortune 1000 companies, the average total cost of downtime is between $1.25 and $2.5 billion per year as it is[9]. Securing against DDoS without diminishing performance requires a careful balancing act.

As mentioned, many companies attempt to mitigate this by redirecting traffic to scrubbing centers that are usually a long distance from the traffic source or origin server, creating a bottleneck that results in latency levels that can be just as bad as an attack. For this reason, limited scrubbing center services are not a realistic choice for blocking DDoS attacks. Furthermore, consider cloud mitigation services with the capability to perform detection and mitigation at locations physically close to an attack source in any global region, for faster response times.

5 CHOOSE SMARTS TO STAY AHEAD OF ATTACKERS

Overcoming increasingly complex DDoS attacks requires more than just a layered approach. It requires you to continuously analyze traffic for malicious patterns that can help you develop the intelligent, adaptive defenses you need to fend off future attacks. The DDoS attack that's underway now is the secret to defeating the next.

When evaluating cloud-based mitigation, it's important to look beyond capacity or transfer and filtering speeds, but also for what kind of intel is enabled by its reach. The larger and more robust the mitigation network, the richer the intelligence it can provide on evolving attack patterns—and the more pre-emptive these services can be.

How Cloudflare can help

Cloudflare's layered security approach combines multiple DDoS mitigation capabilities into one service that prevents disruptions caused by malicious traffic while allowing good traffic through—keeping websites, applications, APIs, and entire networks up and running with high availability and performance.

With data centers in 200+ cities and over 90 countries, and over 35 Tbps network capacity, Cloudflare mitigates DDoS attacks close to the source, within 100 milliseconds of 99% of the Internet-connected population of the developed world.

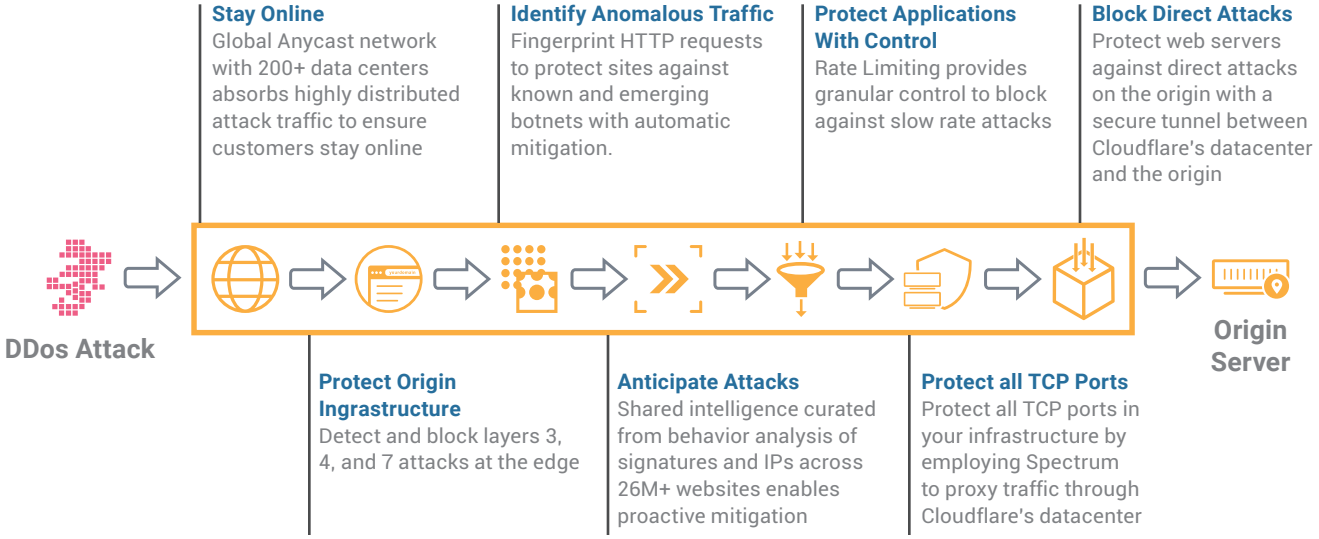

FAST, AUTOMATED MITIGATION

Unlike traditional solutions with bottleneck dependency on limited scrubbing centers, our points of presence globally host security services to protect against DDoS attacks of any size or complexity. That includes the ability to scatter traffic across distributed servers to the point where it is absorbed by the network.

THREAT INTELLIGENCE AT GLOBAL SCALE

Cloudflare's DDoS protection is fueled by the intelligence of its global network, which protects over 25 million+ websites and has over 1 billion unique IPs passing through it every day. This intelligence enables a unique vantage point to protect against the most sophisticated attacks.

COST-EFFECTIVE PROTECTION

All Cloudflare plans offer unlimited and unmetered mitigation of DDoS attacks, regardless of size, at no extra cost—and no penalty for attack-related spikes in network traffic.

EASE OF USE AND MANAGEMENT

Cloudflare's always-on, cloud-based DDoS protection is built on an intuitive interface that empowers users to quickly and easily protect their Internet properties against DDoS attacks of any scale or sophistication in just a few clicks.

INTEGRATED SECURITY AND PERFORMANCE

Our protection is designed to integrate, learn, and operate seamlessly with other security and performance solutions, including Web Application Firewall, Bot Management, Magic Transit, Load Balancer, CDN and more.

DATA ANALYSIS, YOUR WAY

Cloudflare Analytics enables you to analyze DDoS events through Cloudflare's integrated dashboard or GraphQL. Alternately, Cloudflare logs can be integrated with leading third-party SIEMs to seamlessly integrate with your existing business processes.

Conclusion

An effective strategy for meeting the challenges associated with DDoS attacks requires a comprehensive approach that addresses all threats at all layers. While on-premise solutions can be part of the answer, they can quickly get expensive. A more robust solution will integrate performance with scalable, cloud-based mitigation that provisions services at the network edge for maximum agility and unlimited capacity, ensuring resiliency against DDoS attacks of any size or complexity.

Footnotes

1 Kaspersky Labs, “DDoS Breach Costs Rise to Over $2M for Enterprises Finds Kaspersky Lab Report," Kaspersky Labs, February 22, 2018

2 Crane, Casey, “The 15 Top DDoS Statistics You Should Know in 2020," Cybercrime Magazine, November 16, 2019

3 DeNisco Rayome, Alison, “Major DDoS attacks increased 967% this year," TechRepublic, April 24, 2019

4 Avital, Nadav, “2019 Global DDoS Threat Landscape Report," Security Boulevard, February 5, 2020

5 Cook, Sam, “DDoS attack statistics and facts for 2018-2019," Comparitech, August 20, 2019

6 Ibid

7 Stein, Jake, “Behind the Buzzword: The Reality of Real Time," InformationWeek, September 5, 2019

8 Tyson, Mark, “Users Lose a Full Working Week Every Year Due to Slow Computers," Hexus.net, October 2013

9 “IDC Study - The cost of downtime," Tech Republic, September 30, 2017

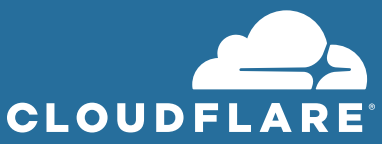

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

© 2020 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.

REV: 200330