CLOUDFLARE

WHITEPAPER

Understanding the role of cloud-delivered network protection

Content

Understanding the role of cloud-delivered network protection

Every organization has public-facing network infrastructure. This infrastructure hosts a range of applications, including those that support employees, customers, and business partners. It may include gateway services that provide network connectivity such as VPNs, virtual desktop servers, or jump servers. This infrastructure also delivers critical IT and network services such as email, file servers, DNS, remote access, and communications like VoIP.

Public-facing infrastructure is a challenge to secure, as it is addressable by anyone on the Internet. As a consequence, it is vulnerable to a number of different threat vectors. For example, attackers can scan the infrastructure to inventory discoverable applications and services. If attackers find a vulnerability in the operating system, the appliance, or the software in the future, they could take advantage of the exploitable window before a patch is developed or installed.

Attacks can also take other forms as well. For example, distributed denial-of-service (DDoS) attacks are a longstanding threat to public-facing infrastructure that can directly cause financial losses and brand damage. These attacks can easily overwhelm the threshold of conventional protections, and security teams learn about the issues only after the attack is already underway.

In light of these network challenges, organizations have deployed a long list of firewall helpers, such as additional network appliances as well as filtering services. The additional functionality comes at a cost, as many make the networking inefficient, and the protections are still ineffective against the growing sophistication of the threat landscape.

Conventional network firewalls, along with multiple firewall helpers, are simply not enough. The time for architectural change is now, but the path for change has not always been clear. In this paper, we will illustrate the types of attacks that are employed against public-facing infrastructure, highlight the problems organizations face, and make the case for implementing cloud-delivered protections through a connectivity cloud.

Challenges with protecting public-facing infrastructure

Security and networking teams use network firewalls to establish the demarcation between the trusted network (internal), the public-facing network or demilitarized zone (DMZ), and the Internet. This is a perfectly valid role for a firewall. However, public-facing infrastructure is subject to a number of attacks that cannot be stopped by firewalls.

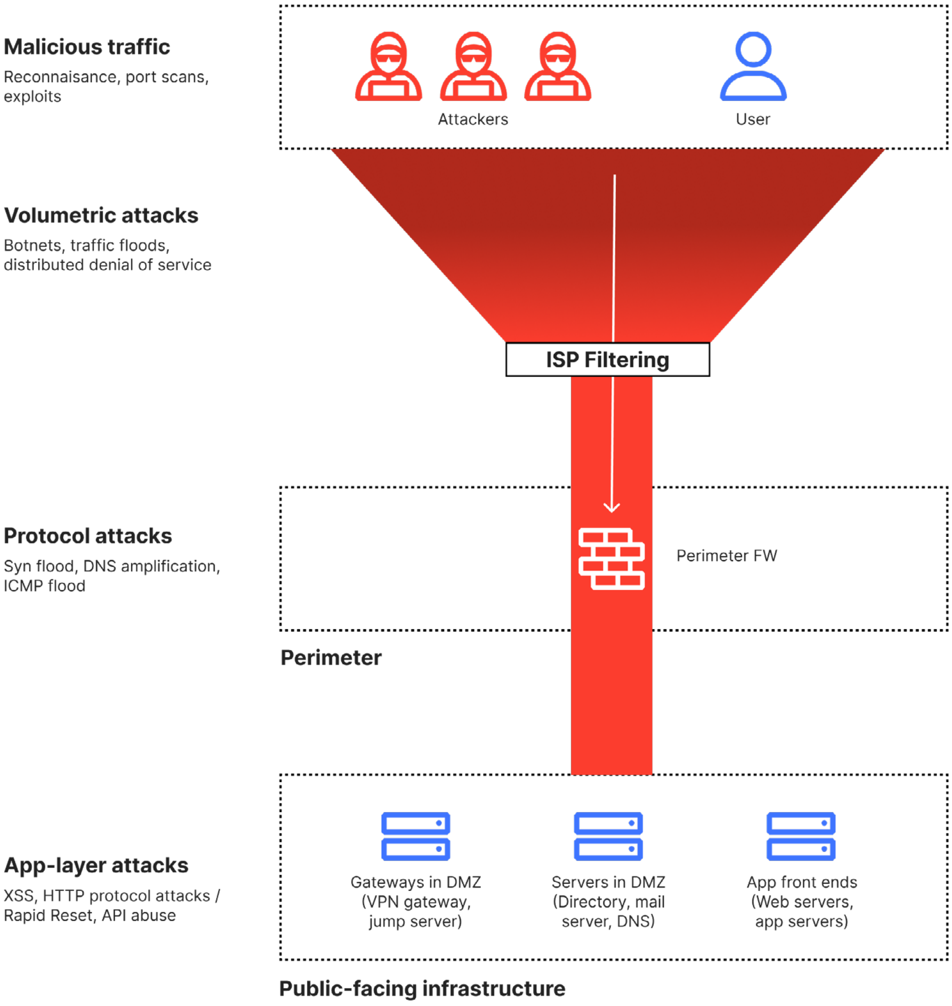
Malicious traffic
Reconnaisance, port scans, exploits
Attackers
User
Volumetric attacks
Botnets, traffic floods, distributed denial of service
ISP Filtering
Protocol attacks
Syn flood, DNS amplification, ICMP flood
Perimeter FW
Perimeter
App-layer attacks
XSS, HTTP protocol attacks / Rapid Reset, API abuse
Gateways in DMZ (VPN gateway, jump server)
Servers in DMZ (Directory, mail server, DNS)
App front ends (Web servers, app servers)
Public-facing infrastructure

A traditional public-facing network architecture.

DDoS: Firewalls provide filtering capabilities, and many do provide protection against types of DDoS attacks. However, firewalls cannot always protect against variants in DDoS techniques. As a result, the effectiveness of their built-in protections fall far short of what's needed for today.

Volumetric DDoS attacks: Volumetric attacks (sometimes called L3 DDoS attacks) deliver an overwhelming amount of traffic to process, measured in gigabits per second (Gbps). To deliver that traffic, these attacks typically use botnets with the assistance of an amplification vulnerability. With sufficient traffic, a volumetric attack saturates the line and overwhelms the firewall's interface. To address these threats, organizations employ a firewall helper, such as upstream filtering with traffic scrubbing centers or ISP filtering. But both introduce their own sets of problems, such as additional latency. ISP filtering has upper limits of what it can process as well, which means it drops what it can't handle and blocks access to legitimate users.

Protocol DDoS attacks: Protocol-based DDoS attacks (sometimes called L3/L4 DDoS attacks) attempt to overwhelm the number of sessions through techniques such as SYN floods. Protocol attacks are measured in size by packets per second (pps). Modern firewalls do have the ability to recognize a protocol attack, but firewalls can only drop protocol attacks if they have sufficient resources to recognize, analyze, and drop the attack that's already underway. Firewall vendors sometimes use these capabilities to support their claims that they offer DDoS protection, without clarifying that protocol attacks can still overwhelm the firewall and being opaque about the other types of DDoS attacks that it cannot stop. To address a firewall's shortcomings, organizations sometimes employ scrubbing centers or ISP filtering as helpers to drop traffic upstream of the firewall.

App-layer DDoS attacks: App-layer DDoS attacks (sometimes called L7 DDoS attacks) differ from both volumetric and protocol-based attacks in that they use communication with the app-layer protocols to cause the denial of service. As such, they are measured in requests per second (rps). Since this type of attack operates from the app layer, it can succeed using a smaller number of hosts. Firewalls (even firewalls with L7 inspection) do not understand the interaction between a host and server, and thus require yet another set of firewall helpers, such as web application firewall (WAF) and DDoS appliances designed to filter threats from interacting with public-facing infrastructure.

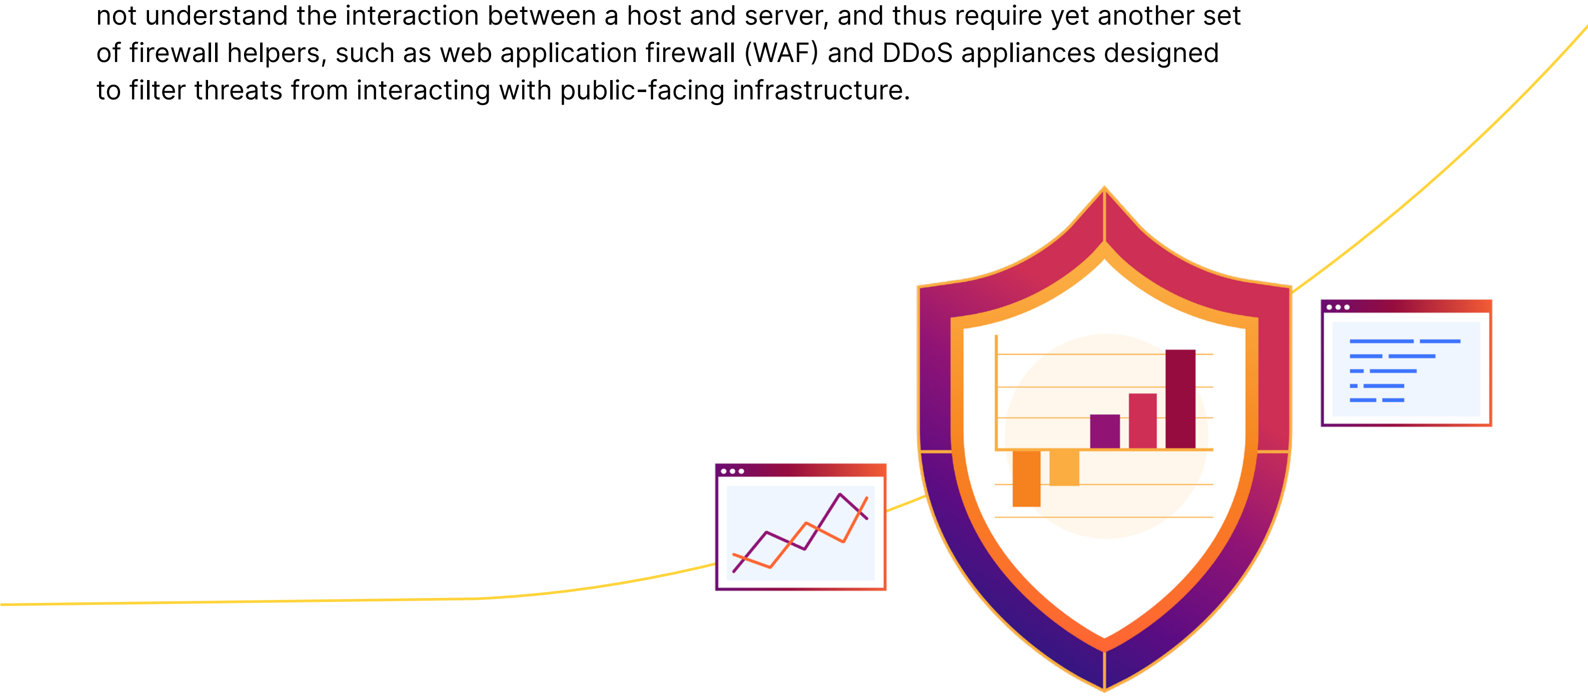

Exploitation of HTTP interfaces: Both applications and network infrastructure frequently provide web interfaces. For example, a VPN appliance has web interfaces for functions such as clientless VPN, which is publicly exposed. These applications are subject to pre-authenticated attacks that exploit a vulnerability, and thus organizations must employ protections to filter and block application abuse.

Scaling firewall protections: Firewalls have finite capacity and are inelastic. When organizations buy a firewall, they must buy a model with sufficient capacity. Their ability to scale firewall protections is constrained by that purchasing decision. Scaling is often a problem because the most cost-efficient decision is to buy only the capacity you need. But when organizations buy insufficient capacity for the lifespan of the firewall, they might be unable to operate computationally heavy protections such as TLS decryption. It's costly to purchase additional capacity to handle atypical traffic volumes. Therefore, capacity is almost always a challenge: there is no scenario that's just right, and many that are suboptimal.

			Firewall Helper Controls					
Types of attacks	**Examples**	**Firewall**	**IP S/IDS**	**WAF**	**API Security**	**DDoS Appliance**	**ISP Filtering**	**Scrubbing Center**
Malicious traffic	Reconnaissance, port scans, expoits	✓	✓					
Volumetric (L3) DDo$ Attacks	Bonets, traffic floods, distributed denial of service						✓	✓
Protocol (L3/ L4) DDo$ Attacks	Syn flood, DNS aplification, ICMP flood	✓					✓	✓
App Layer Attacks	App Abuse: XSS, Bot, SQL injection			✓				
	App-Layer (L7) DDoS: HTTP protocol attacks / Rapid Reset					✓		
	API abuse: Expoiting vulnerabilities, MITM, Credential stuffing				✓			✓

Organizations use "helper controls" to address functionality that firewalls do not perform.

The problem with firewall helpers

To deliver supplemental protections, organizations employ a number of firewall helpers, such as specific app-layer protections (including WAF and DDoS filtering), scrubbing centers, and ISP filtering. These protections exist largely because the organization's needs extend beyond what its firewall can do, and yet the protections introduce problems in their own right.

Additional on-premises appliances: In the past, the only way to add more network protection was to add more network appliances. Inserting an appliance in the network is no simple matter, as the network effectively has to be re-engineered for every new capability the organization needs. Appliance insertion for inline protections causes network downtime, redundancy, and the tug of war choosing between networking uptime and security risk exposure when deciding to fail closed or fail open.

Scrubbing centers: A scrubbing center promises to filter traffic before it reaches the organization. However, the traffic diversion through a scrubbing center introduces a performance hit, since the location of a vendor's scrubbing center is not likely to be a 1:1 match with the customer's destination infrastructure. In the worst cases, it can lead to inefficiencies such as the "trombone effect," which causes traffic to take an undesirable path and damages the overall user experience. In effect, the architecture of scrubbing centers amplifies the philosophical differences between security and networking as the protection comes at a cost to the optimal network paths.

ISP filtering: ISP filtering promises to provide a more efficient networking path for upstream filtering than a scrubbing center. That's because the ISP sits in front of the organization's own public-facing infrastructure. However, while ISP filtering provides more capacity than the organization's firewall, ISP filtering has its own limitations. ISPs are absorbing all of the attack's traffic, which makes the filter highly vulnerable both to the upper limits of volumetric and protocol attacks. What's worse is that ISP filtering is only looking at specific types of attacks and thus requires yet another firewall helper to address the unmet needs.

Defining requirements for architectural change

The very core of the problem is that it's extremely difficult and expensive to absorb a dynamic attack with static resources. In addition, backstopping the firewall with helpers creates its own architectural and management problems. What's necessary is an architecture that can address the functional needs and global scale of the threat landscape with dynamic protection.

The cloud provides a model for elastic compute and networking, but not all clouds are the same. A cloud network designed to deliver security within the networking is necessary. The path forward is a "connectivity cloud," one that is capable of providing global networking to connect any user or office to any application. A connectivity cloud can also deliver a set of composable services designed to recognize and neutralize threat activity.

To address network protections, the necessary security services must include:

Global coverage: With a growing customer base and a rising number of hybrid workers, businesses must be able to provide protections across larger and larger geographies. As such, there are pressures to deliver broader coverage than what's conventionally feasible with an on-premises strategy.

Perimeter firewall services: To protect public-facing infrastructure, a connectivity cloud must provide perimeter firewall services designed to enforce inbound traffic inspections and policy enforcement.

DDoS protections: The connectivity cloud should also offer single-pass protections against volumetric, protocol, and app-layer DDoS attacks. Any model other than single-pass enforcement introduces unwanted latency.

Observability: The connectivity cloud must provide visibility into the current state, behavior, and performance of the network, which allows for better troubleshooting, analysis, and optimization.

Cloudflare's connectivity cloud

Cloudflare's connectivity cloud can help organizations modernize their network and better protect public-facing infrastructure. Built on a composable, programmable architecture, it provides networking and security services across cloud-enabled business infrastructure and applications. As a result, the connectivity cloud can address both current and future needs in your modernization journey.

With Cloudflare, your organization can easily add functionality and support new use cases by enabling services rather than inserting appliances. With a connection to a Cloudflare Anycast data center, you can configure and deploy services from the unified management interface to process traffic. You can address current needs while putting a platform in place to accommodate future use cases and support your entire network modernization journey.

Instead of building out global infrastructure, you can benefit from the lightning-fast performance of the Cloudflare global network. With direct connections to nearly every service provider and cloud provider, the Cloudflare network is within 50ms of 95% of the world's Internet population.

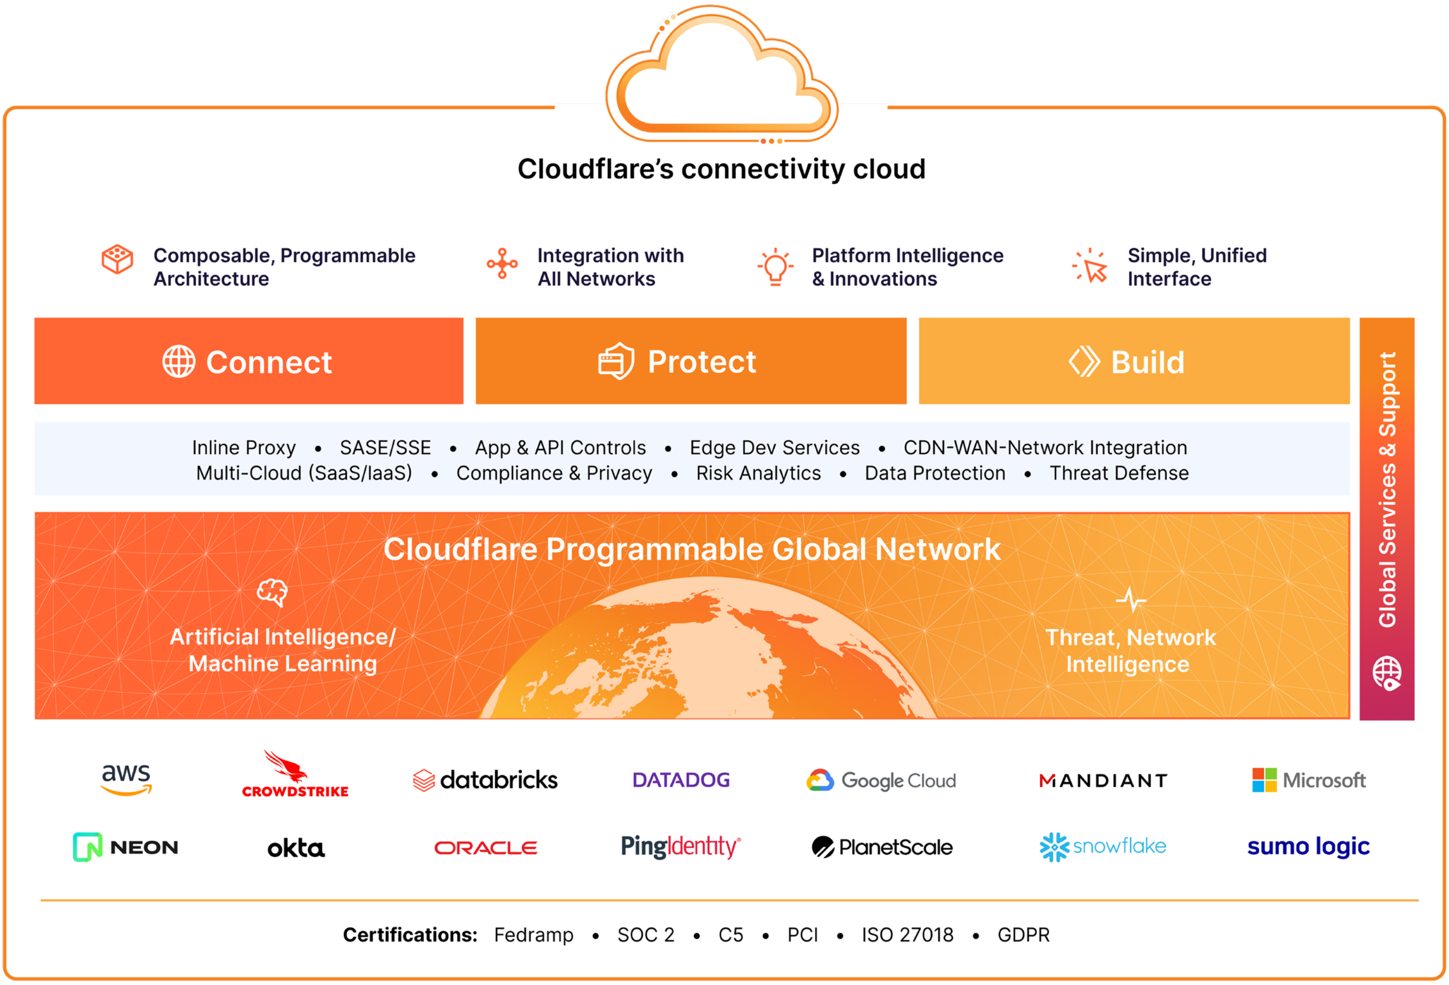
Cloudflare's connectivity cloud
Composable, Programmable Architecture
Integration with All Networks
Platform Intelligence & Innovations
Simple, Unified Interface
Connect
Protect
Build
Inline Proxy • SASE/SSE • App & API Controls • Edge Dev Services • CDN-WAN-Network Integration
Multi-Cloud (SaaS/IaaS) • Compliance & Privacy • Risk Analytics • Data Protection • Threat Defense
Cloudflare Programmable Global Network
Artificial Intelligence/ Machine Learning
Threat, Network Intelligence
Global Services & Support
aws
CROWDSTRIKE
databricks
DATADOG
Google Cloud
MANDIANT
Microsoft
NEON
okta
ORACLE
PingIdentity
PlanetScale
snowflake
sumo logic
Certifications: Fedramp • SOC 2 • C5 • PCI • ISO 27018 • GDPR

The Cloudflare connectivity cloud can help organizations connect, protect, and build everywhere.

How Cloudflare delivers network protection

Cloudflare uses a combination of connectivity cloud services to deliver network protection. To protect public-facing infrastructure, the connectivity cloud blends services together in network locations around the world.

Cloudflare connectivity cloud

Volumetric Attacks

- Absorbs volumetric attacks by diffusing traffic across Cloudflare's network

Network protection

- Firewall enforcement with global policy updates in seconds
- DDoS filtering for volumetric and protocol attacks

After: Cloudflare network protection

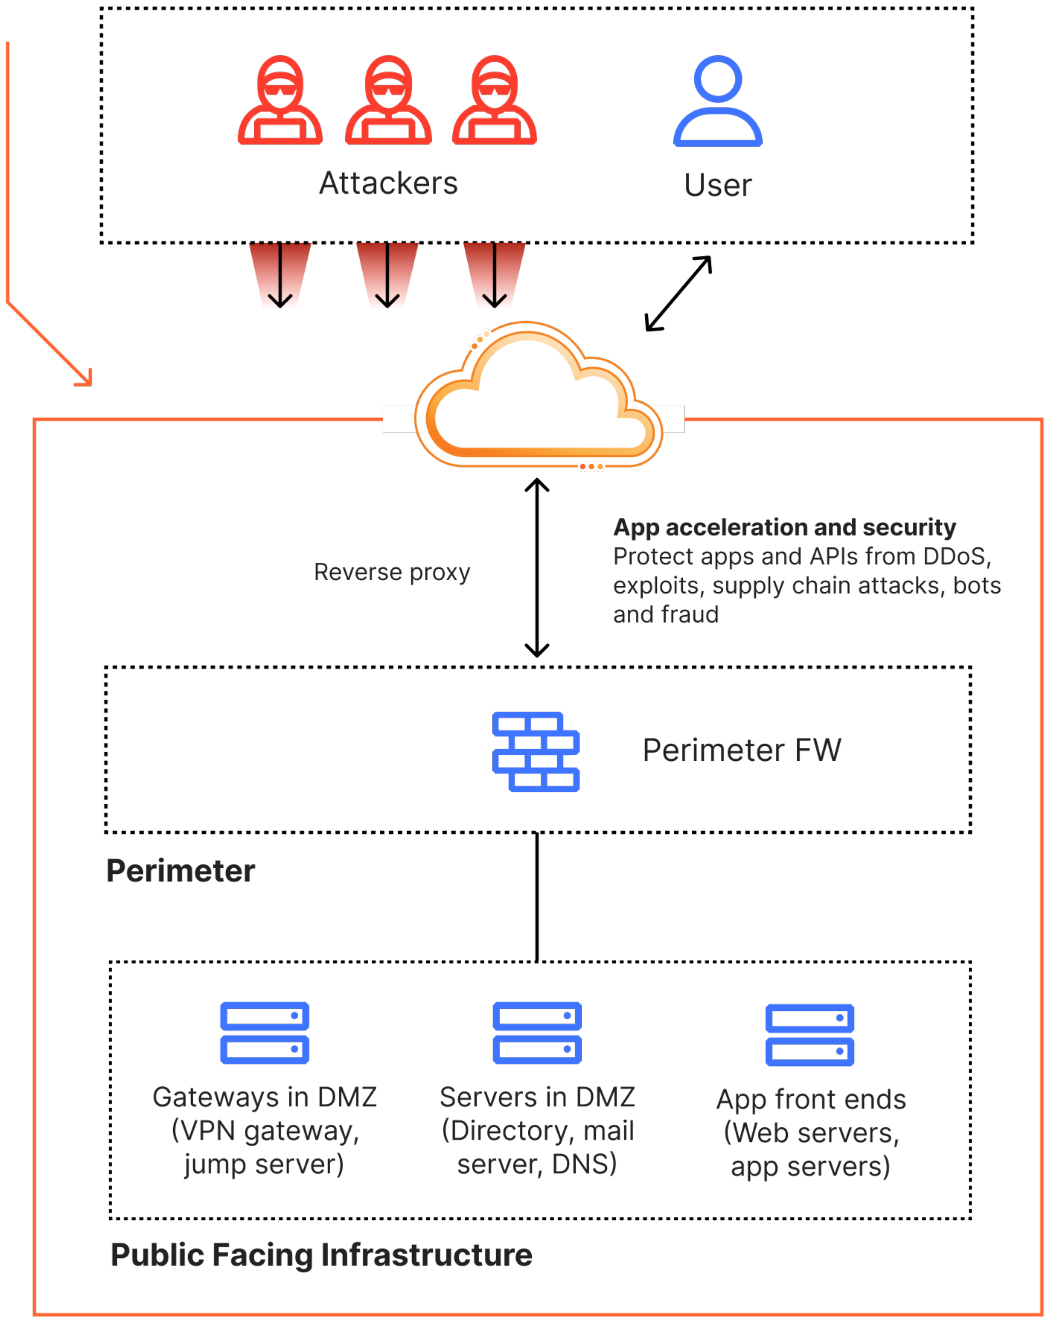
Attackers
User
Reverse proxy
App acceleration and security
Protect apps and APIs from DDoS, exploits, supply chain attacks, bots and fraud
Perimeter FW
Perimeter
Gateways in DMZ (VPN gateway, jump server)
Servers in DMZ (Directory, mail server, DNS)
App front ends (Web servers, app servers)
Public Facing Infrastructure

The Cloudflare connectivity cloud offers a new model for protecting public-facing infrastructure.

To deliver inline protections, all inbound traffic to public-facing infrastructure routes through a Cloudflare data center first. Cloudflare offers a number of ways to connect to the Cloudflare network. One of the underlying concepts behind many of the traffic steering methods is a concept called Anycast. The Cloudflare Anycast network (which covers 330+ cities and counting) acts as the front door to an organization's public-facing infrastructure.

Every Cloudflare data center operates at massive scale with interchangeable compute resources. These data centers operate far above the capacity of what any single organization might build out for themselves. With the Anycast network, all of this capacity is amplified, because all of the Cloudflare data centers help to filter and diffuse attacks.

Consider, for example, how botnets deliver malicious traffic. Botnets use massive numbers of distributed, compromised hosts to generate overwhelming amounts of traffic, packets, or requests. The traffic targets the same destination. With the Anycast network, participants of the botnet see the closest Cloudflare data center that lies in the path toward its intended victim's public-facing infrastructure. The scale of the Cloudflare network "diffuses" traffic volume while simultaneously applying DDoS protection at L3, L4, and L7. In effect, Cloudflare uses the resilience of the distributed network as a countermeasure against an attack.

Cloudflare Magic Firewall provides the front line of inspections that help organizations define what is or is not allowed. It removes the garbage that your existing infrastructure does not need to process, using elastic resources to free up on-premises firewall's limited static resources for other functions. It helps organizations use a defense-in-depth strategy to efficiently unlock value that was otherwise unobtainable due to high on-premises firewall utilization.

Because DDoS attacks affect different layers, organizations must implement protections at each one. Cloudflare's connectivity cloud provides multi-layer protections through the following services:

- **Cloudflare Magic Transit:** Delivers DDoS protection against volumetric attacks at L3.
- **Cloudflare Spectrum:** Provides DDoS protection against protocol-based attacks at L3/ L4.
- **Cloudflare DDoS:** Protects against app-layer DDoS attacks at L7.

These protections work together, and they are all delivered and managed from the same Cloudflare control plane. They eliminate the need for firewall helpers and complement an organization's defense-in-depth strategy, helping to ensure that organizations can maintain operations in light of today's threat landscape.

The Cloudflare connectivity cloud also provides a number of additional services for operating public-facing infrastructure, such as:

- **Cloudflare Network Interconnect** for directly connecting your public-facing infrastructure to the Cloudflare network. Think of Cloudflare as both the on and off ramp for your traffic.
- **App acceleration** for optimizing the delivery performance of your applications.
- **App security** for implementing protections such as a web application firewall (WAF) for traffic passing through Cloudflare.

Next Steps

This defense-in-depth architecture using the Cloudflare connectivity cloud is easy to deploy — far easier than the operational steps for traditional appliance insertion. Once you establish the traffic path to Cloudflare, protections can be enabled with policies deployed globally within minutes.

To learn more about Magic Transit, please visit the reference architecture.

This document is for informational purposes only and is the property of Cloudflare. This document does not create any commitments or assurances from Cloudflare or its affiliates to you. You are responsible for making your own independent assessment of the information in this document. The information in this document is subject to change and does not purport to be all inclusive or to contain all the information that you may need. The responsibilities and liabilities of Cloudflare to its customers are controlled by separate agreements, and this document is not part of, nor does it modify, any agreement between Cloudflare and its customers. Cloudflare services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied.

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

REV:BDES-6349.2024AUG30