DISTRIBUTED SECURITY SOLUTION BRIEF

EQUINIX

CONTROL FUNCTION SERIES DISTRIBUTED DENIAL OF SERVICE (DDOS) PROTECTION

DDoS Protection in an Interconnection Oriented Architecture®

Executive overview

Enterprises are on a cloud transformation journey that requires shifting infrastructure from a traditional, siloed and fixed state to one that is agile and dynamic. The IOA® Playbook offers a compelling strategy for how to do this with Equinix Performance Hub® on Platform Equinix®. Security functions, including Distributed Denial of Service (DDoS) detection and mitigation, can be distributed to solve complexity and integration challenges, as seen in the Distributed Security Digital Edge Playbook. In conjunction with segmenting the network, DDoS protection is one of the first security control functions enterprises must design and implement to prepare their infrastructure for digital business.

DDoS growth has created additional needs

Digital business depends on uninterrupted service. However, DDoS attacks on cloud applications, websites and other services are on the rise, nearly tripling from 2017 to 2018.[1] Attacks against third-party data centers and cloud services increased threefold YoY.[2] Hackers launch DDoS attacks for a number of reasons, including theft of intellectual property, political motivation, financial gain and more. Enterprises undergoing cloud migration are particularly vulnerable, as the movement of workloads to the cloud creates additional, and highly exploitable, attack surfaces that are irresistible to bad actors. As a result, enterprises are grappling with how to defend themselves. The trouble is, patchwork systems and approaches mean enterprises often end up deploying multiple solutions in their efforts to control the situation.

While next-generation firewalls and intrusion prevention systems (IPS) offer some narrow DDoS mitigation, these solutions cannot combat the myriad of possible DDoS attack vectors. In addition, they are overwhelmed by large/complex attacks and can be bypassed by clever attackers. Enterprises have a number of options for architecting and deploying DDoS mitigation. Let's explore these in more depth.

On-premises DDoS protection solutions

Traditionally, enterprises have deployed on-premises, bespoke solutions, usually in a physical or virtual form factor that lives in the data center. On-premises solutions have the advantage of low latency, as the appliance is located in the data center adjacent to application servers. This is especially true if the DDoS protection appliance and the application servers are on the same platform. Some on-premises appliances can also be deployed on demand and activated during an attack, ensuring latency is not added during peacetime. Reduced detection latency and increased quality of detection are additional advantages of on-premises solutions. These devices can examine the entire packet and perform extensive matching to determine bad versus good traffic. Lastly, on-premises-based solutions provide a greater degree of perceived control.

1 Information Age, "Cyber Crime on the Rise: DDoS Attack Volumes Have Trebled in Past Year, Says Study," February 8, 2019,
2 NETSCOUT, Press Release "NETSCOUT Releases 14th Annual Worldwide Infrastructure Security Report," March 20, 2019,

However, an on-premises DDoS protection solution is not without its problems. Available budget is always a consideration. The cost of an on-premises DDoS detection and mitigation appliance can range from tens to hundreds of thousands of dollars, depending on the size and scale required—a cost that is usually paid in full before benefits are received. Additional costs for support and maintenance, as well as resources and staff to manage and maintain the equipment, must also be added (consuming either internal resources or bundled as part of cloud mitigation services costs). Finally, as DDoS attacks increase in size, on-premises DDoS protection solutions may be unable to deter large-scale attacks, due to network capacity issues. Therefore, enterprises may require a subscription to cloud mitigation services as a backup.

Cloud-native solutions

Organizations with some form of distributed architecture and/or services hosted within a cloud service provider often consider a cloud-native DDoS protection solution. These solutions offer significant amounts of available scrubbing capacity and can handle volumetric attacks that on-premises solutions can't. Cloud-native DDoS protection solutions can stop DDoS attacks before they reach the enterprise network or cloud-hosted assets, and they're available either on demand or always on.

An enterprise deploying a cloud-native solution can reap the following advantages:

- **Protection for cloud-based services.** On-premises DDoS protection solutions cannot protect cloud-based applications.
- **Additional scrubbing capacity.** A cloud-native solution can mitigate volumetric DDoS attacks that overwhelm transit link capacity.
- **Lower-cost, subscription-based pricing models.** On-premises DDoS protection solutions can run from tens to hundreds of thousands of dollars, often required to be paid in full at the outset. Cloud-native DDoS solutions enable enterprises to expand or contract DDoS mitigation as needed, while minimizing the size of an on-premises appliance or eliminating it altogether.
- **No added or unnecessary latency.** An on-demand solution incurs no additional latency when an enterprise is not under attack, as peacetime traffic is not diverted to the cloud-native solution.
- **Less management overhead.** A cloud-native DDoS protection solution does not require the same amount of administrative resources.

However, cloud-native solutions also have drawbacks. For example, an on-demand, cloud-native DDoS protection solution does not provide full-time protection and may experience significant delays in detecting and mitigating attacks. On-demand cloud DDoS protection solutions detect DDoS attacks using metadata, which means only configured static traffic thresholds, statistical baselining and very limited signature matching (L1 through L4) can be utilized. As a result, alerts will appear only after incoming traffic exceeds thresholds for a number of minutes, during which time the offending traffic has been analyzed and deemed to be excessive/anomalous. This, of course, delays the application of mitigation efforts.

Once an attack is detected, traffic must be diverted using Border Gateway Protocol (BGP) to the cloud-native solution. It takes some time for the detection, mitigation configuration and forwarding to occur, compromising user experience or worse. Once traffic has been diverted, all inbound traffic flows through the DDoS service provider network, adding latency. The amount of latency will depend primarily on the location of the scrubbing centers and the distances from the destination, although nearly all use anycast addressing and routing methodologies at scrubbing centers to ensure that traffic processing is as geographically optimized as possible. Furthermore, while a cloud-native solution is convenient, this convenience is paid for with a commensurate loss of control.

Hybrid DDoS protection solutions

A hybrid DDoS protection model combines the on-premises appliance model with the notion of expandable cloud capacity to protect against DDoS attacks. In a hybrid DDoS protection approach, traffic flows directly to the on-premises part of the solution. The device inspects incoming traffic for DDoS attacks. When it detects an attack that is too large for it to handle, the appliance signals a cloud-based scrubbing center to perform routine cloud mitigation which includes traffic diversion, scrubbing and reinjection back to the enterprise. Upon completion of the attack, the cloud operator ends the mitigation by withdrawing the BGP prefix, allowing inbound traffic to re-route back to the normal path to the enterprise.

As we've seen with the other deployment models, a hybrid approach offers advantages and disadvantages. Even so, many security analysts identify the hybrid approach as a best practice. Advantages of a hybrid solution include:

- **Low latency.** The on-premises component of the solution enables low latency, by virtue of its location in the data center. Volumetric attacks requiring cloud mitigation will change the path of inbound traffic, adding some latency. However, when the on-premises component and the cloud-native part of the solution are located in proximity, cleaned traffic is returned to the enterprise more efficiently.
- **Real-time and greater fidelity attack detection.** As traffic flows through the on-premises solution at all times, attacks can be detected and dealt with immediately and with greater visibility due to full-packet (L1 through L7) vs. metadata inspection (L1 through L4).
- **Flexible scrubbing capacity.** Enterprises are able to mitigate any size of attack.
- **Shared roles.** The on-premises appliance will continue to perform real-time deep detection and protection during a cloud mitigation, which often augments cloud scrubbing by picking up changes in attack vectors and other indicators the cloud service might be unable to detect (L1-L7 vs. L1-L4).
- **Coordination.** During a coordinated hybrid mitigation, some vendor solutions will allow blacklists to be forwarded from on-premises to cloud, which aids in initiating and updating the cloud mitigation inspection.
- **Control.** The enterprise retains a level of control with the on-premises solution component.

There are few disadvantages to the hybrid model. Since a hybrid solution is comprised of both an on-premises appliance and a cloud-native service, the cost can be higher than that of a cloud-native service only. In addition, management overhead must be factored in to synchronize the on-premises and cloud-native components.

DDoS Protection in an Interconnection Oriented Architecture (IOA)

An enterprise with cloud-hosted applications and services as well as existing data centers running services should consider a hybrid DDoS protection solution to defend against this sort of cyberwarfare. Both cloud-hosted applications and services housed within existing data centers must be protected. Distributing the architecture to multiple, distributed interconnection hubs or exchange points in a virtualized fabric provides traffic visibility in both locations, enabling the detection and mitigation of all potential DDoS attacks via a single solution.

A distributed hybrid DDoS protection solution can therefore be located between the data center and the cloud service provider, in physical proximity to both locations. This physical proximity minimizes latency of returned cleaned traffic to services in either location that require rapid (i.e., real-time) response. In addition, an IOA strategy reduces backhaul costs for cleaned traffic and reduces or completely eliminates GRE/IPsec tunnels typically used to return cleaned traffic to the enterprise.

The image below depicts a hybrid multicloud architecture with DDoS protection services applied.

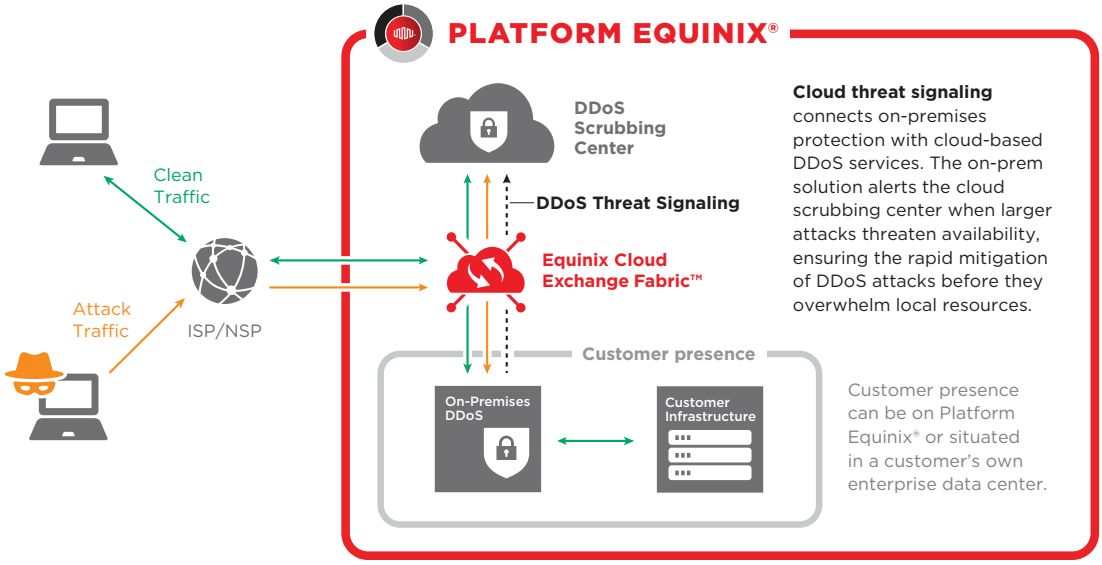

Conclusion

By deploying a hybrid DDoS protection solution within network hubs on a distributed platform, an enterprise can detect and mitigate DDoS attacks targeting on-premises and cloud applications. A virtualized interconnection fabric offers a hybrid DDoS protection solution with visibility into all traffic, enabling detection and mitigation of all potential DDoS attacks via a single solution.

Where to get help

Platform Equinix offers a number of on-premises, cloud-native and hybrid DDoS protection solutions. For more information on how DDoS protection applies in your environment, please contact securityteam@equinix.com.

SB_DDoS_Control_Function_US-EN | 6196 | v090519 |0619 | Q418