radware

On-Demand, Always-on, or Hybrid?

Choosing an Optimal Solution for DDoS Protection

Whitepaper

SHARE THIS WHITEPAPER

Table of Contents

Introduction

Distributed denial of service (DDoS) attacks have caused severe service interruptions and financial damages to organizations throughout 2015. Radware's 2015-2016 Global Application and Network Security Report revealed that over 50% of organizations have experienced some type of DDoS attack in 2015. Yet, as much as 50% of the organizations cited that they are unprepared for such attacks. DDoS attacks are increasing in quantity and severity as these attacks become increasingly complex and persistent. Typical DDoS attacks have evolved to include simultaneous multiple attack vectors that test simple mitigation techniques. Attacks using dynamic IP attacks that challenge mitigation through simple blacklisting are now ubiquitous. Volumetric network-level DDoS attacks at staggering throughput rates of hundreds of Gbps and hundreds of millions of packets per seconds have become commonplace, disabling organizations' network and infrastructure. SSL-based and application-level DDoS attacks that are effective in exploiting bottlenecks in the IT architecture of enterprises have become more prevalent.

Fortunately, there are good solutions to address the threat of DDoS attacks. Solutions include DDoS protection appliances installed on-premise as well as cloud-based DDoS protection services that can be consumed either on-demand or via always-on deployments. Another alternative for DDoS protection is a hybrid approach which combines on-premise DDoS protection appliances and cloud DDoS protection services to provide a robust protection suite. Each of these approaches for deploying DDoS protection has its own benefits, but also bears some challenges. The most appropriate approach for the deployment of DDoS protection depends on the organization's IT architecture and business needs.

On-Premise DDoS Protection Appliances

DDoS protection appliances are powerful technologies that mitigate DDoS attacks. Installed on-premise in the organizations' data center, the best of these appliances detect and mitigate DDoS attacks at all layers, including network-layer, SSL-based and application-layer DDoS attacks. Alternative implementations include a single appliance for both the detection and mitigation of DDoS attacks, and distributed implementations in which detection and mitigation are done by different components and are combined to provide protection from DDoS attacks. Using DDoS protection appliances on-premise has several benefits, as the time it takes to detect and mitigate DDoS attacks is usually minimal compared to other approaches. Since the organization's inbound traffic is not diverted or routed through a cloud DDoS protection service, minimal latency is added in peacetime or during an attack. In addition, when using on-premise appliances that include SSL-based DDoS protection, there is no need to share the organization's certificates with a third party. Also, by handling all traffic with on-premise appliances, the organization can avoid potential regulatory challenges associated with sharing its traffic with a third-party service provider such as Privacy Acts and PCI-DSS certification.

Unfortunately, there is one thing that on-premise, DDoS protection appliances cannot do: provide protection against massive volumetric DDoS attacks that saturate the internet pipe. Massive volumetric DDoS attacks use throughput rates of hundreds of Gbps and hundreds of millions of packets per seconds to overwhelm upstream networking gear, rendering any downstream appliance installed on-premises. Moreover, on-premise DDoS appliances can protect applications in the enterprise data center, but are ineffective in protecting applications hosted on public cloud infrastructures, which have become increasingly ubiquitous as organizations migrate to the cloud. In addition, some organizations have limited ability to install DDoS protection appliances in their data centers or lack the in-house expertise to configure and manage these devices.

Cloud-based DDoS protection services allow enterprises to overcome these challenges. Using cloud-based scrubbing centers strategically deployed worldwide and interconnected for global load balancing, these cloud-based DDoS protection services can absorb volumetric DDoS attacks several orders of magnitude larger than any organization is capable of handling. The alternative deployment architectures for consuming cloud-based DDoS protection services can be broadly categorized into:

- On-demand cloud-based DDoS protection service in which the organization's incoming traffic is diverted to cloud-based scrubbing centers only upon the detection of volumetric attack that threatens to saturate the Internet pipe.

- Always-on cloud-based DDoS protection service in which an organization's traffic is always routed through a local point-of-presence (POP), thereby allowing the cloud DDoS protection service to continuously monitor the traffic and mitigate any DDoS attack.

- Hybrid cloud-based DDoS protection service in which on-premise DDoS protection appliances mitigate most attacks locally and are coupled with a cloud DDoS protection service for volumetric attack protection.

On-Demand Cloud-Based DDoS Protection Services

In on-demand cloud-based DDoS protection services, the detection of DDoS attacks is usually done via the remote monitoring of the internet link utilization by collecting flow statistics or router SNMP data on periodic basis, usually every few seconds. Upon the breach of a certain threshold (commonly 70% utilization of the link capacity), the cloud DDoS protection service initiates a diversion of the inbound traffic to the nearest cloud scrubbing center where attack vectors are detected and mitigated so that only legitimate traffic returns to the organization. The merits of on-demand cloud-based DDoS protection services are its simple deployment, as no on-premise appliance is required, and the fact that there is no induced latency in peacetime as traffic is diverted to the cloud DDoS protection service only upon an attack. For these same reasons, the on-demand approach is usually more economical compared to the always-on approach.

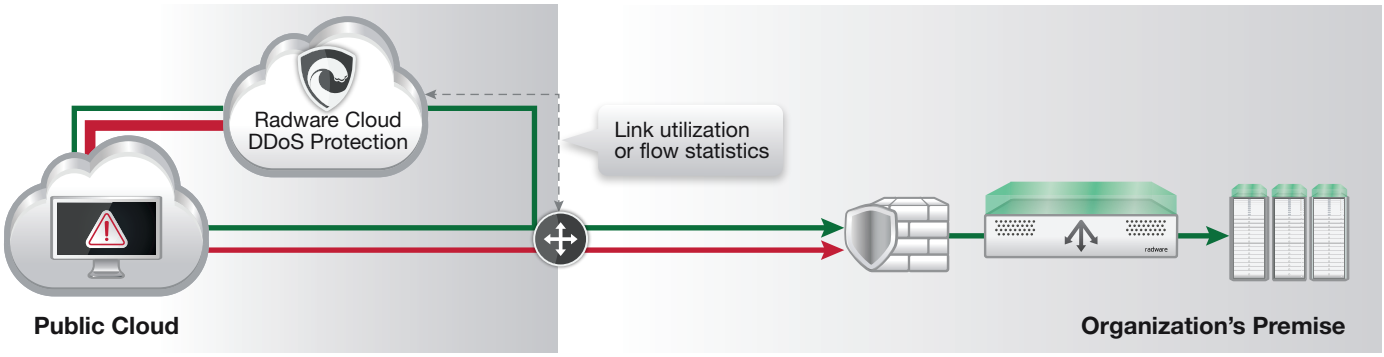

Figure 1: On-Demand Cloud DDoS Protection

On-demand cloud-based DDoS protection services do feature several drawbacks. First, as the detection of DDoS attacks is based on the remote monitoring of the internet link utilization, there is no visibility into any DDoS attack beyond the network layer. Thus SSL-based and application-based DDoS attacks cannot be detected prior to causing severe service disruptions. Moreover, network layer DDoS attacks that congest the link and disrupt the service but remain below the link utilization threshold will usually go undetected. As these type of attacks have grown in popularity across the last year, organizations must consider such attack vectors seriously. Secondly, on-demand cloud DDoS protection is based on diverting the traffic to the cloud service upon a DDoS attack, usually based on DNS or BGP diversion techniques. Unfortunately, these diversions always take time, ranging from a few minutes to several hours, during which the on-going DDoS attack may cause severe service disruption to the organization. Since traffic is only handled upon initiation of an attack, there is limited baseline information, thereby extending the time to mitigate up to 10 minutes after the traffic has been diverted to the cloud. In addition, the on-demand approach is ineffective in protecting applications hosted on a public cloud as there is usually no access to link utilization data of the public cloud infrastructure.

Always-On Cloud-Based DDoS Protection Service

These challenges do not exist in always-on cloud-based DDoS protection services. In this deployment alternative, the organization's traffic is always routed through the local PoP of the cloud DDoS protection service, including in peacetime. This allows the cloud service to detect and mitigate all types of DDoS attacks at all layers, including SSL-based and application-layer attacks, before they interrupt the organization's services. The always-on deployment alternative is highly compelling, as it offers a 'hands off' approach for DDoS protection. By opting for always-on cloud DDoS protection services, enterprises fully outsource DDoS attack detection and mitigation to a third-party expert, requiring minimal resources from the enterprise's IT organization. With the always-on approach, there is no need for traffic diversions, minimizing the time it takes from detection to mitigation of DDoS attacks is minimal, and no service interruption is induced. In addition, always-on features the only approach to provide DDoS protections to applications hosted in the cloud.

Figure 2: Always-on Cloud DDoS Protection

Unfortunately, always-on cloud DDoS protection services also feature several key drawbacks. As traffic is always routed through the cloud service, some additional latency is induced, including during peacetime. This can be a critical shortcoming for latency-sensitive services such as real-time transactional applications. Secondly, it's more expensive than the on-demand approach, as the organization's traffic is always handled by the cloud service, including during peacetime.

Hybrid Cloud DDoS Protection Service

Hybrid cloud DDoS protection services, in which on-premise DDoS protection appliances are coupled with a cloud DDoS protection service, allows organizations to enjoy most of the benefits of the various deployment alternatives while avoiding most of their drawbacks. In the hybrid approach, an on-premise DDoS protection appliance detects and mitigates DDoS attacks at all layers, including network-layer, SSL-based and application-layer attacks. In the event of a massive volumetric DDoS attack that saturates the internet link, traffic is routed to the nearest cloud scrubbing center, where attack vectors are detected and mitigated. This hybrid approach provides fastest time to mitigate of most DDoS attacks as DDoS assaults are mitigated on-premise and only volumetric attacks are diverted to the cloud. For the same reason, the hybrid approach allows organizations to enjoy minimal latency during peacetime.

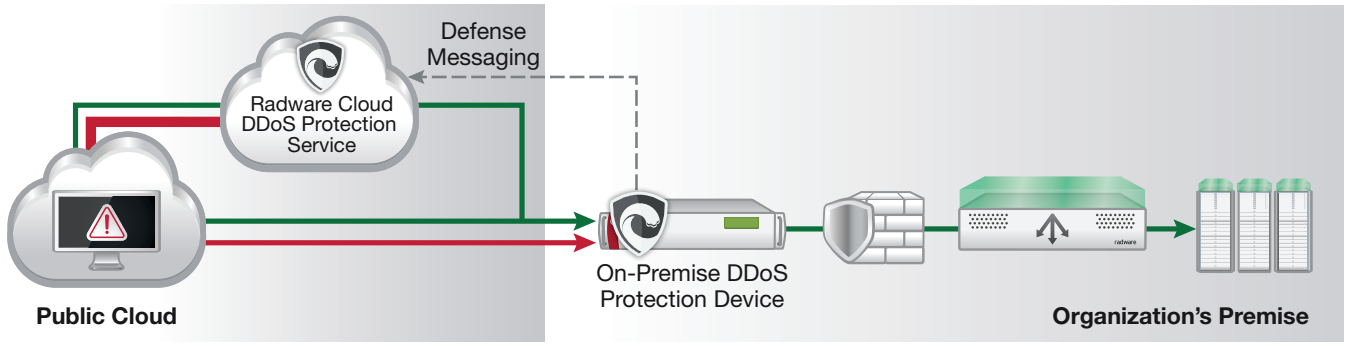

Figure 3: Hybrid Cloud DDoS Protection

Yet, the hybrid approach also features several challenges. First, as the hybrid approach is based on an on-premise DDoS mitigation appliance, it cannot provide an effective DDoS protection to applications hosted in the cloud. Secondly, if the on-premise DDoS solution and the cloud-based DDoS service do not share protection policies and signatures in real time, it can take up to 30 minutes to mitigate a volumetric DDoS attack following diversion to a cloud scrubbing center. This is a common pitfall when the DDoS protection appliance and the cloud scrubbing service are provided by different vendors. To avoid this, when deciding on a hybrid DDoS protection solution, its important to chose one that includes synchronization and messaging of traffic baselines and attack information between the cloud and on-premise components of the solution.

What's the best fit for my organization?

The most appropriate approach depends on the organization's IT architecture and business needs. Several questions should be answered prior to choosing the optimal solution:

- Are the assets that require protection hosted on-premise, in the cloud, or across both via a hybrid deployment model?
- Does the organization have the capacity and expertise to install, configure and manage an on-premise DDoS protection appliance?
- What is the level of sensitivity of the different enterprise services to additional latency during peacetime?
- How sensitive is the organization to SSL-based and application-level attacks, beyond network-layer attacks?
- How sensitive is the organization to the service disruption that may be induced during diversions?

In general, the hybrid approach is the best fit for organizations that have applications on-premise and have the capacity and expertise to handle on-premise appliances. In this case, the hybrid approach provides the fastest time to mitigate most DDoS attacks and the lowest induced latency in peacetime. However, to minimize time to mitigate volumetric attacks after diversions, it would be best to choose on-premise DDoS protection appliances and the cloud DDoS protection service that share traffic protection policies and signatures in real time. This means implementing both solutions from the same vendor. Also, the hybrid approach must be complemented by an always-on cloud DDoS protection service to any applications the organization has that are hosted in the cloud.

The always-on approach is the best fit, and in fact the only solution, for protecting applications that are hosted in the cloud. It is best fit for organizations that lack in-house resources and expertise to handle DDoS threats and seek peace of mind by fully outsourcing DDoS protection services to an expert organization. The always-on approach would also be the recommended solution for organizations that are continuously attacked by large volumetric attacks. In this case, the always-on approach avoids the service disruptions induced by on-going traffic diversions that are required by the other approaches.

The on-demand approach is typically the most economical one. It is a good fit for organizations that have applications on-premise, are less concerned about SSL-based and application-level DDoS attacks, and are less sensitive to the time it takes to mitigate large volumetric attacks. However, this approach requires other means for detecting SSL-based and application-layer attacks. It is advised not to disregard these threats, as such DDoS attacks are becoming ubiquitous in the rapidly evolving threat landscape.

To create the ideal DDoS protection solution, organizations are advised to consider deploying a combination of these approaches. One common combination includes hybrid cloud DDoS protection to protect the organization's data centers, coupled with an always-on cloud DDoS protection service to protect cloud-based applications.

Radware's Cloud DDoS Protection Services

Radware provides a full suite of cloud DDoS protection services that can be deployed in either Hybrid, Always-On or On-Demand cloud DDoS protection services. Organizations can opt to implement one of these deployment alternatives, or choose a combination and benefit from:

- Radware's battle-proven Emergency Response Team (ERT) for on-premise and cloud-based deployments.
- Global network of scrubbing centers with over 2Tbps mitigation capacity and Cloud DDoS Protection Services that are built to detect and mitigate all types of DDoS attacks.
- Market-leading DDoS mitigation appliances, featuring the only cloud DDoS protection service that can automatically generate protections for zero-day attacks within seconds.
- A unique patent-protect technology for mitigating SSL-based attacks, Cloud DDoS Protection Services maintains user data confidentiality and removes the operational dependencies between service provider and the organizations when keys are changed.
- DefenseMessaging, a signaling mechanism that shares protection policies and signatures between Radware's DDoS protection appliances and Radware's cloud security nodes in real time, minimizing mitigation times of DDoS attacks upon diverting traffic to the cloud.

About Radware

Radware® (NASDAQ: RDWR), is a global leader of application delivery and cyber security solutions for virtual, cloud and software defined data centers. Its award-winning solutions portfolio delivers service level assurance for business-critical applications, while maximizing IT efficiency. Radware's solutions empower more than 10,000 enterprise and carrier customers worldwide to adapt to market challenges quickly, maintain business continuity and achieve maximum productivity while keeping costs down. For more information, please visit www.radware.com.

Radware encourages you to join our community and follow us on: Facebook, Google+, LinkedIn, Radware Blog, SlideShare, Twitter, YouTube, Radware Connect app for iPhone® and our security center DDoSWarriors.com that provides a comprehensive analysis on DDoS attack tools, trends and threats.