WHITE PAPER

Hybrid Approach to DDoS Mitigation

corero
FIRST LINE OF DEFENSE

Executive Summary

As organizations consider options for DDoS mitigation, it is important to realize that the optimal solution is a hybrid approach which derives benefits of both an on-premises solution and a cloud-based mitigation service. At a high level, an on-premises appliance provides always on protection against DDoS attacks, in particular application layer attacks, and other unwanted traffic. To ensure maximum protection, an on-premises device can be backed up by a cloud-based mitigation service to defeat sporadic high bandwidth DDoS attacks that exceed the capacity of an organization's Internet pipe. The main benefit of a hybrid approach over just relying on an on-demand cloud-based mitigation is that the on-premises device dramatically reduces the number of times an organization needs to switchover to cloud-based mitigation. This not only lowers the costs and saves time associated with those switchovers but also provides organizations with "always on" protection against all forms of DDoS attacks and other unwanted traffic.

Challenges with Cloud-Based Mitigation

Consider this analogy. Would you visit the doctor each time you sneeze? Most people take some basic precautions at home before deciding to make the trip to the doctor's office. Such a trip would be time consuming and costly in cases where you don't need it. Let's apply the same logic to using cloud-based DDoS mitigation EACH TIME you come under attack. Figure 1 describes how it works.

Step 1: You come under attack and need to determine if it is a real attack and whether it warrants a switchover to cloud-based mitigation. If so, you contact the cloud provider.

Step 2: The cloud provider makes network level changes (DNS redirection, BGP re-routing) to divert your traffic to the cloud provider's scrubbing center.

Step 3: The scrubbing center REDUCES the attack traffic and passes the rest of the traffic (including some threats – described later) to you. Sometime in the future, the cloud-based protection is turned off and operations return to normal.

Challenges

- Does the attack warrant a switchover?
- How long will the switchover take?
- How long will the business be interrupted?
- How much will it cost to switchover?
- When would normal operations resume?

Figure 1: Normal Operation and Cloud-Based Mitigation

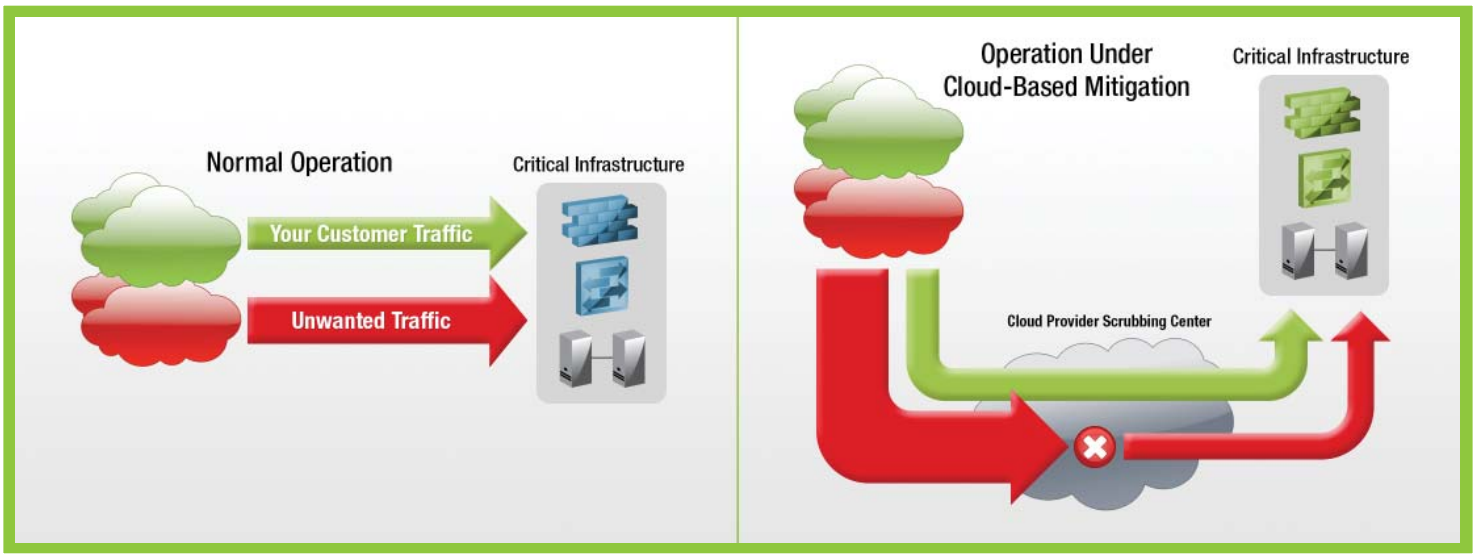

Critical infrastructure of an organization that relies soley on cloud-based mitigation is prone to costly downtime caused by unwanted/attack traffic, resulting in potential revenue losses.

Even the smallest of attacks would require organizations to make a major switchover to cloud-based mitigation, costing many thousands of dollars per switchover.

In cases where attack bandwidth exceeds the size of your Internet pipe, cloud-based DDoS mitigation may be the only option, not to mention the obvious benefit of simply renting DDoS mitigation as needed rather than owning it. However, relying completely on cloud-based mitigation has the following disadvantages:

- Switching over to a cloud-based mitigation is disruptive and can cause temporary outages
- The decision to switchover and divert traffic can take 30 minutes to a couple of hours
- Mitigation costs can vary greatly depending on bandwidth, length of mitigation, and protected sites
- During cloud-based mitigation, you're exposed to threats not mitigated by the cloud provider such as:
  - Outbound communications to botnet command & control servers from botnets inside your organization
  - Low and slow scans and reconnaissance attempts used to profile your infrastructure
  - Any attacks requiring full deep packet inspection which would be too resource intensive to perform on ALL traffic
- Attackers can bypass DNS redirection by directly targeting an organization's IP addresses
- Once traffic returns to normal operation, attacks can re-emerge, restarting the cycle

It is apparent that when under attack, the decision to switchover to a cloud provider only should be taken if absolutely necessary to mitigate the attack. While an organization is contemplating a switchover, it likely is still under attack, rendering the availability of its web presence non-deterministic.

Consider a Hybrid Approach

For the most cost effective round-the-clock protection against DDoS attacks, organizations should consider a hybrid approach, deploying a dedicated on-premises solution and **only** utilizing cloud-based mitigation services for the "full pipe" attacks where attack bandwidth exceeds the organization's Internet pipe. Figure 2 shows how a hybrid approach with Corero's First Line of Defense would provide organizations cost effective, **always on** protection with the option of backing it up with a cloud-based mitigation service for high bandwidth attacks.

Figure 2: Hybrid Approach with Corero's First Line of Defense

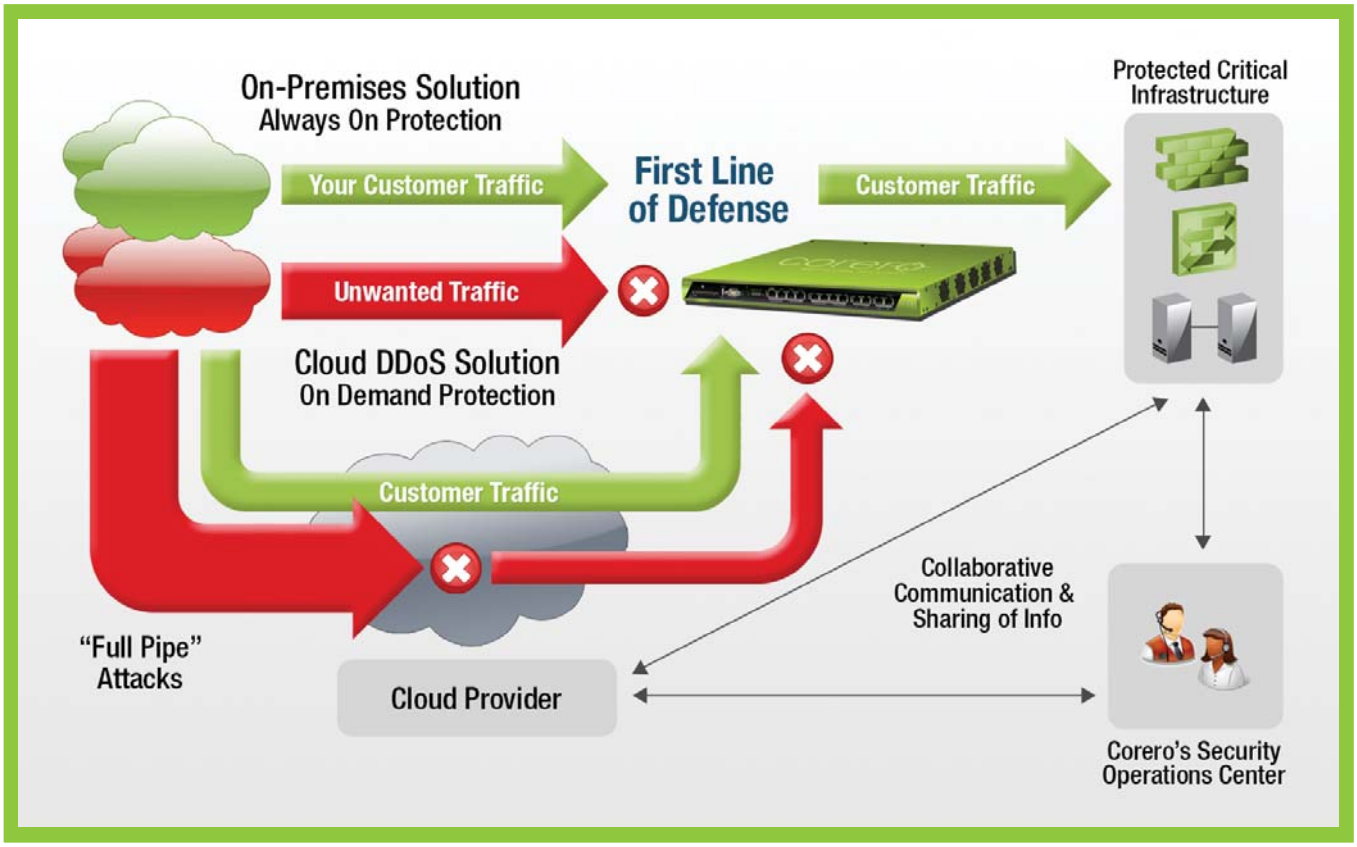

Corero's on-premises DDoS Defense combined with on demand cloud-based mitigation provides comprehensive protection and visibility at the enterprise perimeter.

Deploy an On-Premises Solution for Always On Protection

With Corero's DDoS Defense System (DDS) as an on-premises First Line of Defense solution, your critical infrastructure gets always on protection against DDoS attacks and unwanted traffic such as network level floods, reflective DDoS attacks, outbound DDoS attacks, application layer attacks, specially crafted packets, scans and reconnaissance attempts, and advanced evasion techniques (AETs) . The Corero DDS continuously monitors for deviations in traffic and violations of configured security policies for both **inbound and outbound traffic**. Any unwanted traffic and attack traffic is dropped while legitimate customer traffic is allowed to pass, keeping your organization's critical infrastructure fully operational without suffering costly downtime.

Judiciously Use Cloud-Based Protection

According to a 2012 survey by Neustar Inc., 73% of DDoS attacks are smaller than 100 Mbps, 14% range from 100Mbps-1Gbps, and only 13% exceed 1Gbps. Because more than 80% of the attacks are under 1Gbps, they can be easily handled by an on-premises solution. For large volumetric attacks that exceed the bandwidth of an organization's Internet pipe, a cloud-based solution should be utilized. A switchover to a cloud-based solution in these cases is a significant event, and hence should be done judiciously. There is almost always significant downtime associated with switchover events because of the manual process that requires analysis of the cause of the downtime, human coordination to carry out the switchover and research to begin mitigation/scrubbing. An on-premises solution such as Corero's First Line of Defense can proactively provide advanced alerting of volumetric attacks, attack forensics and attacker IPs in order to make the switchover and mitigation as quick and efficient as possible. Such early warning signs of a larger volumetric attack can be monitored by Corero's 24x7 Security Operations Center (SOC) to help organizations proactively decide whether to consider a switchover to the cloud provider and when to do so, thereby adding additional intelligence and control to their business continuity plan. These proactive measures can not only increase the availability of your organization's critical infrastructure but also save valuable resources which would otherwise be spent on fighting fires.

During the switchover period, an on-premises solution would continue to provide the necessary protection for any threats not mitigated by the cloud-based solution. Continuous monitoring can also provide guidance on when your organization should be switched back to normal operation. Further, collaborative communication and sharing of information among you, Corero, and the cloud provider enables visibility at the enterprise perimeter to help with forensics and compliance as well as enhances the overall performance and security of your network. This type of collaboration can also be used to optimize the level of protection provided by the cloud service provider as well as improve your commercial relationship with them.

Economics of the Hybrid Approach

An important aspect of switching over to cloud-based mitigation is the non-deterministic nature of the cost involved, which ultimately depends on the commercial/contractual agreement with the cloud service provider. Switching over to cloud-based mitigation is a significant event which can result in overage charges if an organization exceeds the number of allowed mitigations. Further, overage charges may be tied to number of events and attack bandwidth peaks, which puts your DDoS budget in the hands of the attacker! Organizations should consider the following questions to help determine how the hybrid approach can save them money:

- If switching over costs 1 hour of downtime, what is the revenue loss to the business?
- What are the costs (mitigation fees and manpower consumed) of an unnecessary switchover?
- How many switchovers would it take to pay for the investment in an on-premises solution?

An hour of downtime can vary across industries but it is not unusual for an on-line retailer or a financial institution to lose hundreds of thousands of dollars in that time. When combining these losses with the mitigation costs incurred of just a single unnecessary switchover to a cloud-based provider, total costs can far exceed the cost of investing in an on-premises solution. A 2012 survey by Neustar Inc. found that organizations in financial services, telecom, and ecommerce can suffer hourly revenue losses from $10K to $100K with a DDoS attack of less than 100 MBps. The chart below (Figure 3) compares the total cost of a single DDoS attack (as short as an hour) with the investment in an on-premises solution.

Figure 3: Investment in an On-premises Solution Could Pay for Itself in Just a Single DDoS Incident

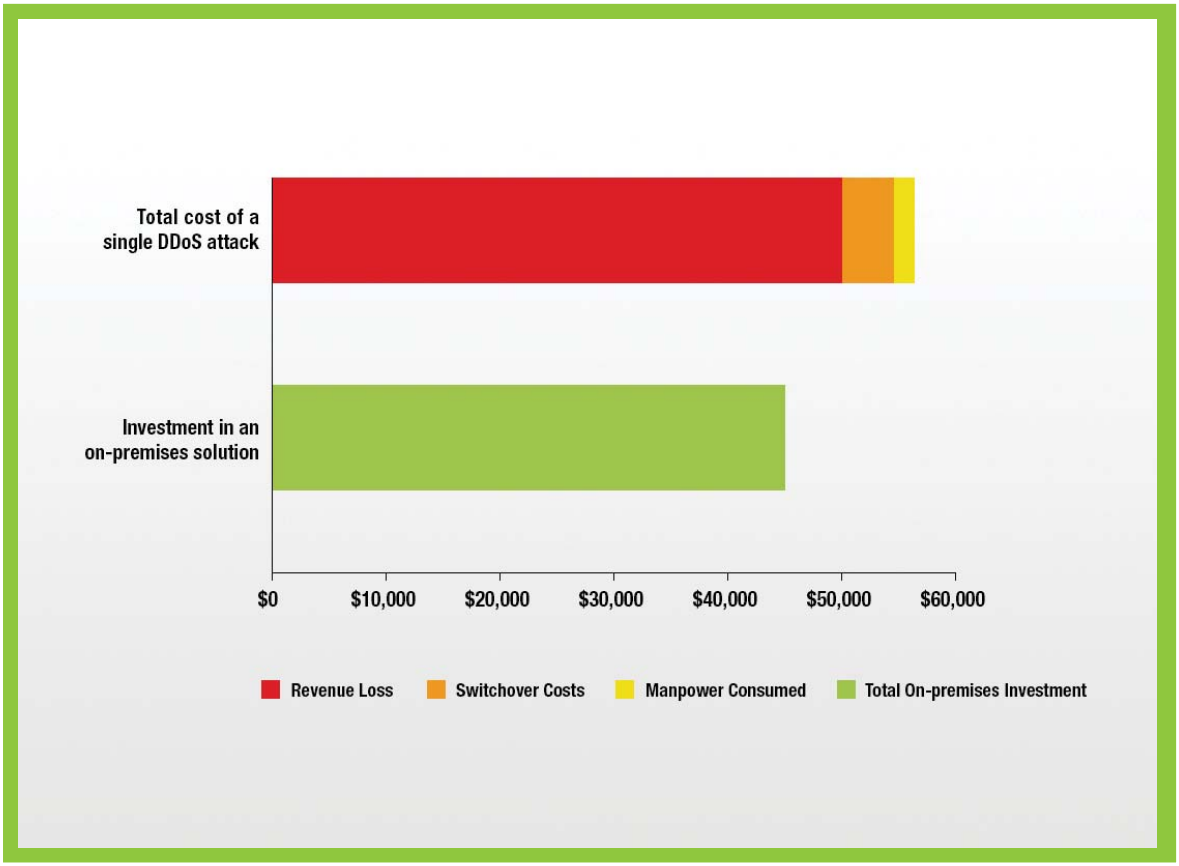

When compared with the total cost of a single DDoS attack, the on-premises component of the hybrid solution could pay for itself in just a single incident. The hybrid approach to mitigation provides your organization with continuous protection against all DDoS variants, enhances your visibility and forensic capabilities and dramatically reduces the number of times your organization needs to switchover to cloud-based mitigation services.

About Corero Network Security

Corero Network Security, an organization's First Line of Defense, is an international network security company and a leading provider of Distributed Denial of Service (DDoS) defense and next generation security solutions. As the First Line of Defense, Corero's products and services stop attacks at the perimeter including DDoS, server targeted, and zero-day attacks, protecting IT infrastructure and eliminating downtime. Customers include enterprises across industries from banking, to financial services, gaming, education, retail and critical infrastructure as well as service providers and government organizations worldwide. Corero's solutions are dynamic and automatically respond to evolving cyber attacks, known and unknown, allowing existing IT infrastructure – such as firewalls which are ineffective at stopping much of today's unwanted traffic at the perimeter – to perform their intended purposes. Corero's products are transparent, highly scalable and feature the lowest latency and highest reliability in the industry. Corero is headquartered in Hudson, Massachusetts with offices around the world. www.corero.com.

Corporate Headquarters
1 Cabot Road
Hudson, MA 01749
Phone: +1.978.212.1500
www.corero.com

EMEA Headquarters
68 King William Street
London, England
EC4N 7DZ
Phone: +44 (0) 207.959.2496