VERISIGN™

WHITE PAPER

APPROACHES TO DDoS PROTECTION:

A COST ANALYSIS

CONTENTS

ALL ORGANIZATIONS WITH AN ONLINE PRESENCE OR DEPENDENCE ON INTERNET-BASED SYSTEMS NEED TO FORTIFY THEIR DEFENSES AGAINST DISTRIBUTED DENIAL OF SERVICE (DDoS) ATTACKS. THESE ATTACKS—WHICH CAN DISRUPT ONLINE SERVICES AND APPLICATIONS—HAVE BEEN CITED AS THE NUMBER-ONE OPERATIONAL SECURITY PROBLEM FACING THE SERVICE PROVIDER COMMUNITY TODAY. IN A RECENT FORRESTER SURVEY, 74 PERCENT OF RESPONDENTS REPORTED EXPERIENCING A DDoS ATTACK IN THE 12-MONTH PERIOD ENDING MARCH 2009.[1]

In 2009, organizations experienced more than 350,000 DDoS attacks,[2] affecting businesses with critical online applications across all vertical and market segments. Attacks are also becoming larger, with peak DDoS attack size growing from about 400 Mbps in 2001 to over 69 Gbps in 2010. Every 26 minutes, on average, an organization comes under an attack larger than 1 Gbps, and every 190 minutes an attack takes place that's larger than 10 Gbps. At the same time, 58 percent of networks surveyed by Forrester reported increased attack complexity and a greater focus on application availability. With botnets available for rent at $100 per day or less,[3] anybody with a grudge or an agenda can easily obtain the computing power they need to launch a DDoS attack, on a scale that will easily overwhelm most organizations.

WHAT DOES DDoS COST AN ORGANIZATION?

DDoS can cost an organization intangible losses and in more subtle ways—whether it has been attacked or whether it is just dedicating resources to defend against an attack. An attack that effectively impacts target systems disrupts the revenue from these systems for the duration of the attack, often with residual effects. Additionally, once the attack is over, the interruption to an organization's service can inflict lasting damage on its reputation, as well as negatively impacting customer satisfaction and trust. Even if an organization has never been attacked, both the far-reaching effects and high probability of DDoS means it will still need to take steps to protect against potential attacks. Keeping up with the DDoS arms race requires spending on bandwidth, hardware, and expertise.

1 Forrester Consulting, The Trends and Changing Landscape of DDoS Threats and Protection, July 2009
2 Arbor Networks The Internet Goes to War, December 2010; Worldwide Infrastructure Security Report, Vol. IV; October 2008
3 www.computerworld.com/s/article/9139398/With_botnets_everywhere_DDoS_attacks_get_cheaper

If an organization's most critical assets are in some way tied to its online availability, then it needs to ensure that availability by protecting against DDoS attacks, both at the network and application layer. Every organization needs to weigh all these costs to select the right DDoS mitigation strategy. It should estimate how much damage a DDoS attack could cause by calculating how much each minute of downtime costs for each of its most critical Internet-facing assets, estimating the probability and projected rate of occurrence of such incidents, and then using those numbers to decide on a budget. Once the organization has decided how much protection it wants and how much it can spend, it is ready to consider the features and costs of the different approaches to DDoS mitigation—traditional on-premises filtering, on-premises or ISP-based intelligent filtering, and cloud-based services—to get the maximum protection from that investment.

DDoS DEFINED

A denial of service (DoS) attack occurs when traffic is sent from one host to another computer with the intent of disrupting an online application or service. A distributed denial-of-service (DDoS) attack occurs when multiple hosts (such as compromised PCs) are leveraged to carry out and amplify an attack. Attackers usually create the denial of service condition by either consuming server bandwidth or by impairing the server itself. Typical targets include Web servers, DNS servers, application servers, routers, firewalls and Internet bandwidth.

ELEMENTS THAT MAKE UP A DDoS SOLUTION

In evaluating the different approaches to DDoS mitigation, organizations should first look at how each of the approaches delivers the three key elements of DDoS defense:

Filtering

While "blackholing" inbound traffic towards a target can help an organization to withstand DDoS attacks and minimize collateral damage to adjacent infrastructure elements and network services, nullrouting, or blocking all traffic would effectively take the target offline—which is exactly what the attackers are trying to achieve. The goal of DDoS mitigation is to filter out only the packets that represent attack traffic, enabling legitimate transactions to continue unimpeded.

Filters operate at different layers of the TCP/IP stack, with the most simplistic of them functioning at the network and transport layers. Simple filters are effective against some attacks, but complex attacks require analysis and filtering based on packet payloads and across multiple transactions at the application layer.

Basic filters provide shallow packet inspection that just checks the packet headers.; however, headers may not provide any indication of what is in the packet, within a given traffic flow, or across multiple traffic flows, so more advanced filters—called intelligent filters—provide dynamic deep packet inspection (DPI) and mitigation capabilities.

The best intelligent filtering mechanisms also use dynamic profiling rather than manual or static lists of attackers when deriving source-based countermeasures. Ongoing analysis of traffic enables these intelligent filters to adapt to ever changing Internet conditions. Over time, state-of-the-art filtering progressively blocks more malicious traffic while preserving business continuity, service availability, and application performance.

Bandwidth and routers

Capacity of bandwidth and routers is often the determining factor in how large an attack an organization can absorb. To withstand a 10 Gbps attack, it will need bandwidth and routers that can handle its normal traffic plus the 10 Gbps of attack

traffic. Additionally, well-established business continuity best practices dictate that critical systems have full-redundancy at both the network and service layer. So, if normal traffic levels are about 7 Gbps, the organization would need an infrastructure that could support approximately 34 Gbps of traffic and application transaction handling capacity.

Expert engineers

In the fight against DDoS attacks, there is no substitute for expertise. Organizations need engineers who understand how to rapidly identify and respond to attacks. Firewall or network experts cannot be as effective in attack conditions as a trained network security specialist, because keeping up with the latest attack trends and defense best practices across network and critical services is a full-time job. Organizations that rely on staff who are not trained to respond to DDoS attacks will find themselves slow to mitigate attacks, frustrated by multi-stage and multivector attacks that can morph at a moment's notice, and distracted from their core security focus, weakening their defenses against other types of malicious attacks. Such organizations might also find themselves confusing malicious activity with flash crowds and other anomalous but legitimate events, leading to an inappropriate incident response.

DEFENDING AGAINST DDoS ATTACKS: FINDING THE APPROACH THAT'S RIGHT FOR EACH ORGANIZATION

In defending against DDoS attacks, organizations typically take one of three approaches: they build their own on-premises traditional firewall router solution, they deploy intelligent filtering infrastructure either in their own data centers or buy it as a service from their ISP, or increasingly they turn to a cloud-based DDoS defense solutions specialist for support.

Scenario 1: Traditional on-premise firewall router protection

On-premise protection using firewalls and routers is the most basic approach, used by most companies. Despite its widespread adoption, this approach has significant filter, bandwidth, and staff limitations that make organizations vulnerable to (and slow to recover from) an increasing array of attacks.

This type of solution uses indiscriminate basic filtering and rate-limiting on the network and transport layers only, with no deep packet inspection that could help flag attacks. Response times are slow because it relies heavily on manual detection, as well as tweaking of policies and tuning of traditional router and firewall security mechanisms in the event of an attack.It also does not scale well, with Access Control List (ACL) size limits and only manual or static filters to block a dynamic list of attack sources and vectors.

Bandwidth and routers are also a point of vulnerability. Organization simply cannot afford to invest heavily in excess capacity that does not add value to the business except to serve as protection in case of an attack. Link capacity is often the bottleneck with on-premise solutions, as attack scale today commonly exceeds even high-speed access link capacities. Likewise, an organization usually

cannot justify a significant investment in expert staff experienced in detecting and mitigating DDoS attacks. Instead, network and firewall experts scramble to solve problems that often span network, firewall and application-layer skill sets. Because these staff might not immediately identify and deploy the right countermeasures, recovering from an attack could easily take days.

SCENARIO 1: TRADITIONAL FIREWALL ROUTER DEFENSE

Characteristics:

- Traditional protection mechanisms only, with no intelligent filtering
- Typically only 10-15 Gbps of protection with little to no protection from application-layer attacks
- In-house firewall or network staff usually lack specialized DDoS expertise

Costs:

- 10 Gbps ISP circuit $10/Mbps and 10 percent minimum commitment = $10,000 monthly recurring charge (MRC)
- $10,000 MRC x 2 for redundancy = $20,000 MRC
- 2x 10 Gbps capable routers >= $200,000 which, at a 36-month depreciation = $5,555/month

Total cost: >$65,000/month/site +$10/Mbps during attacks that exceed minimum commitments

Scenario 2: Intelligent filtering

When organizations first realize that they need more than traditional routers and firewalls to protect them from DDoS-related problems, many either acquire and deploy intelligent filtering facilities themselves, or look to their existing ISPs to provide intelligent filtering as an extension to their current services.

In this approach, all Internet traffic goes first to intelligent filters, which use one or more techniques such as active challenge validation, statistical modeling, rate-based anomaly detection, and deep packet inspection to better protect networks and applications from DDoS traffic while letting legitimate traffic through. Still, organizations might find this filtering solution does not scale well: particularly if they receive Internet traffic through several data centers or ISPs, they will need to set up intelligent filters at each one individually. And many ISPs have only traditional firewalls or a single filtering technology available, which may not be effective against the full spectrum of attack vectors.

While this approach usually provides more resilience, the expense of providing filters and bandwidth for each data center means that this strategy is not effective against the largest attacks. An ISP-based solution usually offers more bandwidth than on-premises intelligent filtering: ISPs can leverage their scale to build up

much greater network capacity, and their upstream filtering eliminates downstream and access circuit congestion riks. However, ISPs can only protect against DDoS attacks that go through their networks, so if an organization also uses multiple ISPs, it will need to implement multiple solutions for complete DDoS protection.

Intelligent filtering solutions typically require dedicated, skilled engineering resources. If a company deploys its own intelligent filtering, it usually does not have dedicated expert staff to help it deal rapidly and skillfully with attack detection and mitigation. While ISPs usually follow standard security incident response procedures—such as implementation of ACLs, real-time blackholing, and ISPspecific features or processes—there is a lot of variation in the level of DDoS expertise ISPs offer.

SCENARIO 2: DEPLOY INTELLIGENT FILTERING PLATFORMS

Characteristics:

- Typically only 10-15 Gbps of protection due to cost of infrastructure
- In-house or ISP staff often lack specialized DDoS expertise
- No cross-ISP correlation or mitigation strategies for detection, reporting, or mitigation

Costs:

- 10 Gbps ISP circuit $10/Mbps and 10 percent minimum commitment = $10,000 MRC
- $10,000 MRC x 2 for redundancy = $20,000 MRC
- 2 x 10 Gbps capable routers >= $200,000 which, at a 36-month depreciation = $5,555/month
- 10 Gbps intelligent filtering platforms cost $100,000 - $200,000, which for a $150,000 platform at a 36-month depreciation = $4,166/month

Total cost: >$29,722/month/site + $10/Mbps during attacks that exceed minimum commitments

Scenario 3: Cloud-based DDoS protection providers

Increasingly, organizations that need to stay online work with a dedicated DDoS solutions service provider. Some of these providers deliver a cloud-based service that is carrier agnostic, enabling the provider to filter all incoming traffic and eliminating the need to deploy multiple solutions for network access service providers. For many, a single DDoS solutions provider service with cross-ISP correlation for detection and mitigation is not only the most effective and easily maintained solution, but also the least expensive, because it enables organizations to avoid significant capital outlay and leverage the DDoS protection service provider's bandwidth and 24x7 access to dedicated experts who can help deal with operational problems.

Organizations should ensure the solutions they consider offer best-in-class filtering, typically through dedicated large-capacity scrubbing centers and surgical filtering with intelligent mitigation devices. Many offer multiple intelligent filtering platforms, some augmented with additional proprietary technology to speed attack detection and countermeasure implementation. Providers may even have established relationships with other networks, in house and external security research organizations, or law enforcement agencies to alert them to attacks and specific new vectors taking place elsewhere, allowing cross-sector correlation and threat intelligence to assist with identifying emerging threats, new attack vectors, and effective countermeasures much more quickly.

Services that operate "in the cloud" provide inherently higher traffic capacity, which means better protection against larger attacks. Internet traffic destined for the target traverses an Internet "scrubbing" center before reaching the target organization's network, preserving availability and application performance for legitimate services. Additionally, a cloud-based DDoS provider can pool bandwidth and infrastructure resources and allocate them as needed, which means a client under attack can tap into an enormous amount of bandwidth and infrastructure investment to absorb much larger attacks than might otherwise be feasible.

Detailed data on baseline network usage can minimize the impact of an attack in two ways. First, it can help experts identify an attack more quickly. Secondly, holistic data on network and application usage can be combined with other data sources to help pinpoint legitimate versus illegitimate traffic, allowing the former to continue through while the latter gets filtered out.

Because DDoS mitigation is a primary focus, these service providers can afford to invest in highly skilled, around-the-clock personnel who make it their business to know everything there is to know about DDoS detection, multi-threaded response, and the most effective countermeasures. These experts work with multiple carriers, clients, networks, vendors, and peers worldwide, gaining the perspective that helps them rapidly distinguish between non-malicious traffic peaks (such as flash crowds) and a traditional or new attack vector. As a result, they can respond more quickly and mitigate attacks more effectively than staff whose main expertise is in networks or firewalls. Some cloud-based service providers also provide monitoring as well as filtering services, watching their customers' traffic for both legitimate and malicious anomalies.

SCENARIO 3: CLOUD-BASED DDoS SERVICES PROVIDER

Characteristics:

- Geographically diverse filtering platforms protect facilities in multiple regions
- More Gbps of total protection capability than most ISPs can provide
- DDoS experts available 24x7

Costs:

- +$10,000/Gbps during attacks > 8 Gbps
- Can include detection technology for up to 4 routers
- Includes 8 Gbps of protection, instantly scalable to more than 350 Gbps
- Additional protection @ $10,000 per Gbps/per >8 Gbps attack
- ~$200,000 annually ~$17,000/month
- Available in a range of protection levels to fit customer risk profiles

Total cost: ~$17,000/month/site

Depending on your organization's needs, it is possible to get more advanced DDoS protection while spending less. The most expensive DDoS protection options—traditional filtering or intelligent filtering, either on-premise or ISP-based—do not offer the same level of protection as cloud-based solutions, even though cloud-based solutions cost much less.

CONCLUSION: RAMP UP PROTECTION WHILE REDUCING COSTS

Expensive on-premise DDoS detection and protection solutions may be adequate for protecting against smaller brute force attacks, but will become increasingly ineffective as DDoS attacks continue to grow in size, frequency, and sophistication. Even the most scalable of these options—an ISP-provided intelligent filtering service—will only protect organizations from a portion of the attack traffic. In contrast, cloud-based managed DDoS mitigation services protect organizations against large-scale attacks across multiple ISPs, providing a single correlation point for detection and mitigation, as well as reporting of attack attributes. In addition to state-of-the-art filtering, specialist cloud-based DDoS protection providers offer the expertise of dedicated engineers who can respond to attacks with proven processes and methodologies, as well as an awareness of the latest developments in new and emerging attack profiles.

It is cost-prohibitive for all but the largest companies to create this level of protection themselves, but highly affordable for them to subscribe to protection services through a worldclass, cloud-based DDoS mitigation solution. And by simply subscribing to a cloud-based solution instead of building and maintaining its own, an organization eliminates the hidden cost of DDoS protection: an array of operational distractions associated with defending against DDoS attacks and continually maintaining purpose-built infrastructure and expertise to mitigate those attacks. Whether an organization's mission is selling sports equipment, providing financial services, managing medical records or providing social networking, chances are it needs internal IT staff to focus on innovations that can help build and retain its position in the market.

Staff simply cannot do that as effectively if they are spending their time and resources just to ensure availability of online services against specific threats to availability. By choosing to outsource DDoS protection to a specialized service provider whose business is DDoS protection, an organization can take advantage of economies of scale, free the creative energy of staff to do what they do best, and empower them with the extra funds that no longer have to go to overprovisioning bandwidth and application serving capacity.

DDoS APPROACHES: COST COMPARISON

	TRADITIONAL FILTERING	INTELLIGENT FILTERING, EITHER ISP OR ON-PREMISES	CLOUD-BASED DDOS SERVICES PROVIDER
Cost per month	$25,556	$29,722	$17,000
Protection included in monthly cost	1-2 Gbps	1-2 Gbps	8 Gbps
Total protection capability	10-15 Gbps	10-15 Gbps	>300 Gbps
Cost of additional protection	$10,000/Gbps	$10,000/Gbps	$10,000/Gbps
Expertise	Existing network and firewall staff	Existing network and firewall staff; possibly some DDoS expertise from ISP	24x7 dedicated security operation center an on-site DDoS expertise

ABOUT THE VERISIGN DDoS PROTECTION SERVICE

The Verisign™ DDoS Protection Service is an in-the-cloud solution that provides scalable, cost-effective protection against distributed denial of service (DDoS) attacks. The Verisign DDoS Protection Service combines people, processes, and technology to offer a massively fortified, comprehensive service that can effectively and efficiently mitigate the world's largest DDoS attacks. The service addresses the facets of DDoS mitigation—from assessment, monitoring, detection, and reporting to DDoS source analysis. The Verisign DDoS Protection Service helps protect organizations from catastrophic DDoS attacks by detecting and filtering malicious traffic upstream of the organization's network with hundreds of customers and partners give Verisign analysts broad visibility into Internet traffic.

This unique insight—along with highly refined layered-filtering technology that works at the network, application, and session layers of the TCP/IP stack— enables the Verisign DDoS Protection Service to quickly and accurately distinguish between malicious and benign traffic, and to forward legitimate traffic to the organization. The military-grade network operations centers are provisioned to absorb the largest DDoS attacks, while proven procedures enable Verisign staff to react quickly and with agility to defend against the rapidly changing threat landscape.

LEARN MORE

For more information, please email Learnmore@verisign.com.

ABOUT VERISIGN

Verisign is the trusted provider of Internet infrastructure services for the digital world. Billions of times each day, companies and consumers rely on our Internet infrastructure to communicate and conduct commerce with confidence.

VERISIGN™

VerisignInc.com