MAZEBOLT

WHITEPAPER

4 Things DDoS Mitigation Vendors Won't Tell You

Table of Contents

DDoS Mitigation Overview

Distributed Denial of Service (DDOS) attackers manage to remotely control computers "bots", otherwise known as systems. They then direct thousands of these systems to send requests to one target, for example a website or application. This causes a traffic flood and leads servers to crash, affecting internal business continuity and preventing legitimate visitors from gaining access. As a result of their power to threaten service performance or to shut down a website entirely, even for a short time, DDoS attacks are a constant threat to businesses.

To ensure DDoS security, enterprises invest in mitigation strategies and technologies to counteract the business risks posed by DDoS attacks. DDoS mitigation solutions resist or mitigate the impact of Distributed Denial of Service (DDoS) attacks on networks by blocking malicious traffic in network traffic and application usage and allowing legitimate traffic to flow through.

How Mitigation Works

DDoS mitigation flows through two essential steps:

1. Detection – is the stage when the solution identifies traffic flow deviations that could indicate an oncoming DDoS attack. This stage is extremely crucial, and the time taken to detect is the goal. Even a few minutes delay could mean that the attack has been launched.
2. Diversion – on detection, cloud mitigation solutions reroute traffic away from their target via DNS (Domain Name System) or BGP (Border Gateway Protocol). On-premises solutions decide whether the traffic should be filtered or discarded. If it is decided to filter out, then DDoS traffic is removed by identifying patterns that indicate that it is DDoS and not legitimate traffic.

Most enterprises continue to rely on various DDoS mitigation postures such as Intrusion Detection systems and Web Application Firewalls. However, these postures may not be fully equipped to save businesses from DDoS attacks. For example, firewalls often create bottlenecks and accelerate outages. More about the advantages and disadvantages of DDoS mitigation postures here.

Vulnerabilities in Existing Mitigation Postures:

Cloud Scrubbing	Content Delivery Network (CDN)	Customer Premises Equipment (CPE)	Intrusion Prevention Systems (IPS)	Web Application Firewalls (WAFs)	Load Balancer	Firewalls
Application layer attacks may be passed over unless the cloud service provider has the relevant decryption keys i.e. "SSL Visibility". and professional services engagement.	CDNs protect organizations against attacks that use the DNS names as their target. If the attacker targets the same organization by inputting the site's IP address the site may be vulnerable.	CPE equipment without a scrubbing center will not protect against large volumetric attacks, or provide protection against internet pipe saturation.	The underlying design is focused on blocking security breaches and is not set to stop a DDoS attack. Thus can only be used to help filter out leakage from components up stream, or potentially to block prolonged Layer 7 attack campaigns.	Proxies, may support undefined HTTP protocol structs for custom applications, but the fact that they are custom usually makes them a vulnerability unless well configured.	Vulnerable to state-table saturation attacks e.g. HTTP attacks and a SYN flood.	The firewall is vulnerable to state-table saturation attacks e.g. HTTP attacks and a SYN flood.
Sophisticated multi-layer attacks require a granular capability for detecting and blocking attacks which scrubbing centers are not always efficient at adapting to.	Advanced attackers can find and attack the source IP of the website directly, circumventing the CDN completely.	Requires manual fine tuning as well as ongoing costs related to infrastructure management	Encrypted packets can remain undiscovered by IPS leaving a gap in security.	Require network re-configuration and tuning as they attempt to "learn" the applications they are protecting. Vulnerable to overloads when perimeter DDoS equipment doesn't mitigate an HTTP based DDoS attack.	Attackers can directly send volumetric DDoS traffic to the custom TCP and UDP communica-tion protocols	Firewalls may be overwhelmed on a perimeter network which is shared with third party platforms.

Innovating to beat the attacks, DDoS defenses have grown in technology expertise and evolved into independent DDoS mitigation solutions that help organizations in mitigating attacks to prevent downtime.

However, the year 2019 saw 8.4 million DDoS attacks. This translates into 670,000 attacks per month or 23000 attacks per day and 16 attacks every minute. The costs of attacks are substantial and a single attack on average cost enterprises US$2 million.

In 2020, the number of attacks soared to 9 million. This means there were 805,000 attacks per month, i.e., 30,000 per day and 21 attacks per minute. The attacks cost enterprises between US$2.3 million to US$4 million. The attacks were more complex and the strategies ever more complicated with multi-vector attacks increasing by 1000% since 2019.

An example of an attack that penetrated despite mitigation is the Dyn attack of 2016. The attacks knocked out access to some high-profile web sites, threw as many packets at Dyn's infrastructure as it could, and the company responded with its own mitigation actions as well as cooperation from upstream internet providers who blocked some of the attack flow. Despite these efforts, it still suffered waves of packets 40 to 50 times higher than normal traffic.
[https://www.csoonline.com/article/3135986/ddos-attack-against-overwhelmed-despite-mitigation-efforts.html]

The year 2020 also saw a 570% increase in bit-and-piece DDoS attacks. But the DDoS highlight of the year was the AWS DDoS attack that lasted 3 days and peaked at a 2.3 terabytes per second.

A recent survey by Kaspersky found that 20% of companies with 50 employees or more reported that they have been the victim of at least one DDoS attack.

The question then arises as to how attacks continue to penetrate despite mitigation solutions.

Here are 4 things that DDoS mitigation vendors won't tell you.

The patterns of DDoS attacks are always changing – The attacks growth, intensity, and penetration time, change fast, along with the Internet world. Even with the most sophisticated DDoS protection systems available today, configuration changes cannot keep up-to-date with the ongoing changes to network and services, nor with the changing nature of attacks (shorter in duration but extremely complex) that leave almost no time to react. Over the years, it has been an established practice for enterprises to evaluate existing DDoS mitigation solutions and opt for one based on pricing, popularity, features, and robustness. However, it appears that this solution may not always work. Even the best of solutions, if not updated to reflect network changes, could pave the way for DDoS vulnerabilities. Based on 420 DDoS tests conducted by MazeBolt on enterprises for the first time between 2015 to the end of 2017, configurations leave an inline vulnerability of >48%, empowering DDoS attack vectors to penetrate the best of DDoS mitigation solutions.

This is because enterprise IT teams will not be aware of the vulnerabilities until attacks strike. In these instances, the IT team along with the mitigation solution will rise to action to mitigate the attack. The time taken to mitigate an attack is dependent on the IT team's agility and expertise to detect traffic anomalies and the speed with which the alert is communicated to internal stakeholders. The mitigation solution if it is an on-demand

solution, will also need the IT team to communicate to the mitigation solution's contact person. The mitigation process to begin is dependent on, as we can see, a set of steps that need to be initiated and kickstarted and the whole process could take time allowing the attack to gather speed and cause disruption to services. If it is an always-on solution, the process will be automatically initiated but we need to bear in mind the fact that the mitigation begins after the attack has already been initiated!

In an ideal world this entire process flows through without a hitch. However, in the real world, this process is set up for failure, and is not foolproof. For example, it is a fact that enterprises must constantly deal with employee churn. New staff could make changes to applications and networks without realizing the impact on the DDoS mitigation solution's effectiveness. The larger the enterprise, the more chances that the impact of network changes on DDoS mitigation may not be communicated and fall through the cracks or as in this case, pave the way for DDoS vulnerabilities.

DDoS mitigation solution vendors will later discuss the prowess with which they were able to mitigate the attack. But the fact remains that the solutions by their inherent weakness left the enterprise open for attacks by not closing vulnerabilities as and when they occurred.

False positives – Mitigation solutions often cannot tell the difference between real users and a DDoS attack. DDoS mitigation solutions can conclude mistakenly, a surge in traffic to be attacks and block the same. The reason for this is, mitigation solutions are attuned to react but they cannot and do not have the time when under what could be an attack, to intelligently differentiate between DDoS and legitimate traffic. The impact of legitimate traffic being blocked need not be spelt out. All enterprises with online presence and those that thrive on online traffic know the impact of blocking legitimate traffic. Research indicates that down downtime can cost up to US$100,000 per hour.

Smokescreen – In the recent past there have been several cases where large service disruptions came in parallel with other attack vectors where DDoS was a smokescreen. Mitigation solutions do not offer full protection from these Dark DDoS attacks as they are referred to, as they rely on human observation and intervention which results in a time lapse. In the interim gap, some damage may already be done, such as affecting a network, or stealing confidential data. As mentioned earlier, DDoS mitigation solutions rise to defend against attacks after they are launched. What if the attack is a distraction? Can mitigation solutions be aware of this and if yes how? The fact is, troubleshooting only happens when systems are brought down by a damaging DDoS attack.

Combo and multi-focal attacks – Many of today's advanced DDoS attacks combine some or all the approaches described above to launch complex DDoS attacks. Attackers use a complex mix of different attack vectors to a variety of targets, making it much more complex for mitigation systems and services to focus on what's going on, and what to block first. This strategy successfully achieves longer downtime before attack detection and mitigation.

DDoS mitigation solutions require trained experts to manage the same. Most internal IT teams do not have or feel the need to have in-house DDoS mitigation experts and they leave the mitigation management in the hands of the solution providers. As a result, the IT team cannot be 100% sure whether their mitigation solution is working or has been disabled.

Findings from our Research

All the points indicate one important fact and that is, mitigation solutions are reactive in nature. Mitigation solutions do not constantly re-configure and fine tune their DDoS mitigation policies. This leaves their ongoing visibility limited and forces them to troubleshoot issues at the very worst possible time, that is, when systems are brought down by a successful DDoS attack. These solutions are all reactive, only closing DDoS vulnerabilities after a successful attack happens.

There needs to be a persistent visibility of attack surface risks, while maintaining service levels intact. Also, there needs to be knowledgebase-assisted vulnerability remediation, with prioritized action plan. RADAR™, is MazeBolt's new patented technology solution and part of the MazeBolt security platform. Working with any mitigation solution installed, RADAR™ offers superior DDoS coverage and automated DDoS protection. RADAR™ simulates over 100 attack vectors with all public facing IPs 24/7, giving real time visibility to all DDoS vulnerabilities with zero downtime.

About MazeBolt

MazeBolt introduces a new standard in DDoS coverage, automatically detecting, analyzing, and prioritizing remediation across the network, doubling coverage, and virtually eliminating DDoS exposure without shutting down organizational operations. MazeBolt's continuous defense supercharges the performance of CISOs as well as the mitigation service provider.

