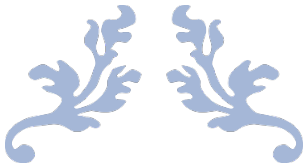

DDOS MITIGATION IN BUSINESS

By: Bryce Green and Mitchell Zuberi

JUNE 30, 2019

Executive Summary

Distributed Denial of Service (DDoS) attacks affect many businesses, both large and small, by controlling zombies to send enormous amounts of packets to the desired system. The system then gets overwhelmed by the influx of traffic and goes down. DDoS attacks are on the rise, both in size and duration. They must be protected against to ensure that key functions of organizations are available to consumers. DDoS mitigation utilizes Detection, Diversion, Filtering, and analysis to combat these attacks. Mitigation methods should be used in conjunction with others to provide the most security and resistance. Some companies such as Metasploit and Distil Networks offer services that are used to fend off attackers.

Uptime is crucial to the success of small businesses, as any downtime could amount to lost sales or traffic. DDoS mitigation can be cost-effective if correct implementation methods are chosen. For small businesses, we believe a combination of ISP protection, third-party services, rate limiting, and IP rotation would consistently protect against almost all threats of DDoS attack. These smaller companies could also utilize data from public honeypots to determine where probable weaknesses are within their systems.

Larger organizations that have more funds to spend on security, can invest in more costly strategies to prevent and analyze the attackers. While we believe they should be utilizing many of the same levels of protection as small businesses, they can also actively create honeypots and use machine learning to better protect against targeted attacks. Machine Learning, in this instance, could be used for traffic monitoring to determine malicious packets from normal traffic. Honeypots similarly deter attacks away from the real network by giving fake access to collect data.

DDoS Mitigation

Online services are a vital part of today's economy. Businesses use web pages to directly interact with customers and perform their day-to-day operations, but this makes them targets for attacks. A common troublesome method performing attacks on these online services is through Distributed Denial-of-Service (DDoS). The goal of this method is to overload the systems that support online services and prevent users from accessing those services, disrupting the operations of businesses and causing financial losses. This goal is reached by sending a high volume of packets or connection requests to a target server in a short period of time that is beyond the capacity the target can handle. What separates DDoS from a normal Denial-of-Service (DoS) attack is that in DDoS the source of the attack is not from a single computer, but from many locations. Attackers often use a network of computers that have been infected and used as bots or *zombies* to coordinate a remote attack. This makes it difficult to stop the traffic from these attacking sources and find the origin of the attack.

Currently, there are no controls that can be put in place to prevent DDoS attacks, but there are ways to protect against them. DDoS Mitigation is the process of protecting an asset from a DDoS attack. While there are a number of methods and techniques that could be used to fulfill the goal of DDoS mitigation, they all fit into at least one of the stages in DDoS mitigation. There are four stages to DDoS mitigation: Detection, Diversion, Filtering, and Analysis (Imperva, 2019) .

Detection is the first stage of DDoS mitigation, where harmful traffic is distinguished from legitimate traffic. Part of DDoS mitigation is the continuity of service without disruption from attacks or from the organization itself. Improperly flagging legitimate traffic as hazardous blocks out potential customers and costs the organization money from that lost customer.

Systems receiving heavy traffic may not necessarily be under attack. For example, sites like *Fandango* were flooded with traffic when tickets for *Avengers: End Game* went on sale. In this case, it is important to look for abnormal traffic that doesn't appear to be eager fans and separate that traffic for the next stage.

Diversion is the second stage of DDoS mitigation and focuses on redirecting potentially harmful traffic away from the systems they were to attack. This process makes the attack more manageable by separating the traffic into chunks that will either be filtered or discarded. It is important to note that this stage can also catch legitimate traffic as well. When this happens, the next stage is responsible for catching that legitimate traffic and passing it onto its' destination. This process is similar to a production line where parts are flagged for defects and separated from the rest to be sent to quality control for verification.

Filtering is the third stage of DDoS mitigation and takes traffic that takes diverted traffic and picks out the DDoS traffic from legitimate traffic. It does this by looking for patterns that distinguish malicious visitors from normal traffic. For instance, the way humans interact with a website differs greatly from the interactions of bots. Utilizing devices that can distinguish between human traffic and unwanted traffic can help in the filtration of packets.

Analysis is the final stage of DDoS mitigation and focuses on learning from information gathered from attacks. Logs should be kept throughout the process of DDoS mitigation and are reviewed to better understand the nature of the attack, attack outcome, and how the correct process of DDoS mitigation can be improved. The more detail there is on an attack, the more thorough the analysis is. This stage is similar to athletic teams reviewing recordings of previous matches to determine what they did well and what they need to improve upon. Their tactics are analyzed and revised to better meet the challenges they face in future matches.

DDoS Mitigation Tools

Many companies provide tools for performing DDoS mitigation, such as *Metasploit* and *Distil Networks*. *Metasploit* is a penetration testing tool developed by *Rapid7*, which finds vulnerabilities in a potential target by performing actions that an attack would take (Rapid7, 2019). This tool helps protect against DDoS attacks by finding vulnerabilities before they can be exploited by an attacker. Since DDoS attacks overwhelm systems with traffic and packets, performing a penetration test with a tool like Metasploit can help determine what that capacity is and whether that capacity should be upgraded to handle future attacks.

Distil Networks provides a Content Protection Network (CPN) to protect systems from bot attacks like those used in DDoS. Traffic is processed through the CPN and analyzed for signs of bot behavior, such as content scraping, and follows user-configured guidelines for responding to those bots. Tools like *Distil Networks* can fight back against the main agents that carry out DDoS attacks before they have a chance to disrupt the services of the target.

DDoS Attacking Trends

Since the development of DDoS attacks, the method has grown in frequency and size exponentially. As reported by *Corero Network Security*, attacks using the DDoS method grew 40% between 2017 and 2018 (Corero,2019). This increase has brought more targets into the crosshairs of attackers, including those without the infrastructure or mitigation tools to shrug off the impact of an attack.

Nothing appears to be off-limits to attacks, as one man from Massachusetts was recently sentenced to 10 years in prison for launching DDoS attacks to bring down services at the *Boston Children's Hospital* and other medical facilities, costing hundreds of thousands of dollars

(Kovacs, 2019). Other public institutions have been targeted, such as universities like the *University of Albany*, which was attacked 17 different times during the month of February this year (*Kupreev*, 2019). While the attacks themselves only brought down services for around five minutes each, the frequency of these attacks made it clear that the University was being targeted with ill intent. Both of these examples were likely carried out by hacktivists, a segment of hackers who perform attacks for political or social reasons. Whatever the reasoning behind these attacks may be, it is clear that any service or organization is at risk.

New technologies and methodologies have also increased in size, scope, and duration of these attacks. The size of attacks are not only determined by the size of packets sent to a target, but also the number of attacking sources. Some attacks this year have been reported as being in the Terabits-per-second range, far exceeding attacks from previous years. *Netscout* has reported that while most attacks remain at around 5Gbps, the maximum attack size has increased by 174% and this will certainly get worse in the years to follow (Netscout, 2019).

The main culprit of these increasing trends is the use of botnets to carry out attacks. These botnets are formed from infected devices from around the globe through the use of malware like the infamous *Mirai* which has caused headaches in recent years. Modern botnets now have the ability to not only infect personal computers but also most household devices. Internet-of-Things (IoT) devices like routers, home appliances, and new smart-home devices like *Nest*, *Alexa*, and *Google* are prime targets for botnet infection due to minimal security measures and large numbers. These devices are in almost every household and business, providing easy entry points for attackers. As such, a DDoS today can be carried out by your refrigerator without your knowledge. Even more troublesome is the fact that many of these tools are publicly

available for anyone to use, enabling *Script Kiddies* to perform attacks on anyone they see fit with relative ease.

DDoS attacks take down services for a few minutes to several hours or longer. Most attacks last less than four hours, but some can last into the tens or even hundreds of hours. Some attacks have lasted nearly two weeks, preventing access to services and costing those business fortunes in lost revenues (Kupreev, 2019). The duration of these attacks between the end of 2018 and early 2019 have already increased and are estimated to continue to do so. Despite the exceptionally high outage times DDoS attacks cause, this aspect is perhaps the most manageable since proper mitigation can filter out malicious packets and restore services even if an attack was successful. This, however, requires that potential targets invest more time and money into developing their defenses.

Traffic Monitoring Through Machine Learning

An attack is only successful if it is allowed to reach its destination. DDoS attacks reach their targets by traversing the internet until they reach the host of a service and deliver their payload, but they can be stopped in their tracks. As discussed earlier, the Filtering stage of DDoS mitigation focuses on finding the attackers before they can impact services. This is done through traffic monitoring, where all incoming traffic is observed for abnormalities. When a police cruiser watches a highway in hiding, they are looking for vehicles going above a certain speed limit. Once a speeder is spotted, the officer pulls them over away from normal traffic before an accident can occur. The same is true for traffic monitoring, where packets or connections that lie outside of the norm are flagged as potential attackers and dealt with by other systems. This technique of DDoS mitigation is a vital step in preventing successful DDoS attacks, but the process may not always catch attackers fast enough before the damage has been done.

Machine Learning is one way of meeting the speed requirements of this technique. A relatively new technology still in its infancy, Artificial Intelligence (A.I.) has the ability to perform some of the basic tasks done by humans at a much faster rate (Atkinson, 2019). A.I. can be taught what is normal traffic and what is abnormal traffic through the process of machine learning. This process is similar to teaching a child the difference between right and wrong, or a new employee on the same task of identifying proper traffic behavior. The difference is that A.I. can handle information on scales humans can't possibly achieve and at levels of accuracy that normal programs simply can't match. Using A.I. to monitor the flow of traffic is ideal for high volume traffic and high volume attacks due to the speed at which an A.I. can detect and filter out connections that are outside of what it has been taught to be normal.

The use of A.I. for monitoring traffic is also a necessary step when confronted with a growing tool for DDoS attackers: A.I. Machine learning is not just in the hands of large organizations but also in some with ill intent. Just like with the benefits of using A.I. to protect against DDoS attacks, A.I. improves the foundations of developing a DDoS attack. If an A.I. can be taught to perform a DDoS attack it is in the best interest of potential targets to leverage A.I. to meet that attack.

One downside to this approach is the fact that the A.I. must be taught what normal is which can lead to false-positives and false-negatives. The problem here is that legitimate traffic may fall outside of the A.I.'s parameters if they are too narrow, but if they are too wide malicious traffic may be let through as well. There are cases, however, where the A.I. may notice a trend among DDoS attacks that had previously gone unnoticed or have been recently developed, in which the A.I. could develop new parameters for detecting and filtering malicious traffic. The algorithms used to train A.I. must be updated over time due to attackers learning

about these algorithms, exploiting their weaknesses, and either working-around the A.I. or fooling it into misinterpreting the nature of DDoS traffic.

Besides its' early stage in development, the main drawback to using A.I. to protect from DDoS or any other form of attacks is the high cost of using one in the first place. While large businesses, organizations, and governments have the resources and funding to support the use of A.I. most entities below that level simply can't justify the cost of such an investment, at least not at this time.

ISP Protection

One solution for DDoS mitigation that would appeal to small businesses is receiving protection from DDoS attacks from their Internet Service Provider (ISP). Traffic to online services must flow through the infrastructure of an ISP before it reaches its destination. ISP's already have infrastructure and expertise in place to handle attacks carried out through their network and can implement DDoS mitigation at a much lower cost than the average business. While the overall effectiveness and response time may match more expensive and complicated strategies, signing up for DDoS mitigation service with an ISP can lower the impact and risk of an attack.

ISP's big and small may use services outside of their own infrastructure, such as *Corero's Smartwall Threat Defense System* (TDS) which is marketed specifically for service providers. The system utilizes *Corero's* security appliances that can scale with the performance required by a service provider with any number of customers. Service providers that use services like those from *Corero* can effectively protect their customers from DDoS attacks without the need for developing their infrastructure.

When an attack is too much for an ISP or their customer to handle, one drastic solution is to implement blackhole routing. Blackhole routing is a last resort solution that can completely disrupt both DDoS and legitimate traffic (Cloudfare, 2019). Traffic to the target is routed to a null pointer and dropped from the network. While this method essentially fulfills an attacker's goal of disrupting traffic to the target a blackhole can be configured with criteria to keep legitimate traffic on the network as much as possible.

Third Party Services

Protection from an ISP is not the only way of defending against DDoS attacks without a target needing to heavily invest in their own infrastructure. Companies like Amazon and Google provide their own services that perform the steps on DDoS mitigation and either stop an attack or reduce the impact one has on a target's services. Amazon's *Route 53* and *AWS Shield* handles DDoS attacks at a scale that most businesses simply cannot match and is hosted at numerous locations across the globe that allows it to handle large scale attacks and be readily available to most locations (Barr, 2018). *Google Cloud Armor* is a similar service that benefits from Google's extensive infrastructure and uses their load balancing services to protect against DDoS attacks scale the capacity limits of incoming traffic with the needs of the business and the size of attacks. Both of these companies have extensive experience with managing and combating DDoS attacks on their own resources and have the infrastructure to extend their services onto organizations without the technical capability or knowhow for a fee. These fees scale with the amount of traffic that is monitored and the number of zones that are protected, with the costs in the range of less than a dollar per million requests or zone.

These services are not going to take care of everything for a business, however, as the business must still take some action to protect itself in case an attack can get around the third-

party service before it can react. In this case, a business should look for cheap solutions to fill the gap so that the overall system of DDoS mitigation is effective. Third party services are themselves at targets for attack due to their high profile and the services they provide. If an attack is successful at slowing down or completely taking down a third-party service for a period of time their clients are left vulnerable to attack.

Rate Limiting

A popular solution in fighting DDoS attacks is to throttle the number of requests your system can accept at a given time. This is done through rate limiting where a system is configured with a ceiling for connection requests from a given source of traffic. An administrator for a system might know that the average number of requests coming from a single user or IP address is typically ten requests per minute, so they set the maximum number of requests per minute to that number (Mahdi, 2017). When a source of traffic exceeds that limit any request above that ceiling is ignored and that source can be flagged for potential malicious activity. This technique is also favorable for limiting traffic to the capacity of the system itself to ensure that it is not overwhelmed. The drawback for this method is that while it can limit the requests that a single source can send, a sufficiently large botnet attack could perform enough requests in mass to defeat the limit placed on individual sources.

Honeypots

Honeypots are decoy networks that mimic potential targets for cyber-attacks. They are commonly used to lure hackers away from legitimate targets and gather data on the attack, such as what they want, how they are gaining access, and what methods they are using to attack the network. These systems are often hosted on virtual machines so that if they are infected or compromised, they can be restarted with minimal downtime. Honeypots are a major aspect of

monitoring and defense, as there is no legitimate reason for an end user to try to gain access to a company network. Therefore, if a user attempts to gain access to a honeypot, their activity will be logged and secured for analysis and the creation of preventative measures.

Honeypots should be utilized by large businesses. While they are fairly expensive, they provide valuable insight into hacker's intentions. These systems can also hinder hackers due to the high visibility of fake vulnerabilities. Hackers often attack the easiest targets presented to themselves and these honeypots protect real damage from taking place. Large businesses need to ensure that they have the best security posture possible, and honeypots add awareness to weaknesses that a system may face.

There is a lot of detail that goes into designing a decoy network because it must appear to be authentic. The four steps that go into creating a honeypot include honeypot environment configuration, structure the logging system, configuration of the firewall, and testing the honeypot (Soenke, 2016). Honeypots can be configured in Windows and Linux environments. These environments can be configured on a virtual machine so that if there is any security damage, it wouldn't truly harm the network. Logging is important for the analysis of information that these breach attempts provide. However, there are many different aspects that can be logged through a honeypot. Application events, login attempts, or even file access attempts. These file access attempts are often crucial because there need to be logs of what files are changed within a system and can be taken through versioning of the files. Some networks may require traffic monitoring so tools like Wireshark could be used to produce those logs through packet detection. The next step in the creation of a honeypot is firewall configuration. The firewall is necessary to ensure that only necessary ports should be used to create the honeypot. No method of entry to the real internal network should take place. These ports must be selectively chosen so that hackers

believe that there actually is a vulnerability and it isn't too good to be true. The last step in the creation of a honeypot is testing. Testing is important to ensure that there are no blind spots in the system logs/monitors. This could include penetration testing, as well as port scanning and utilization of an IDS.

In DDoS mitigation, honeypots are typically accessible through a DMZ which makes them accessible. It will lure an attacker into believing that the system has been compromised and gather data as necessary for the organization. All other traffic that is not malicious should be forwarded to the destination that was intended. The real difficulty and cost of running a honeypot is the illusion. For example, all responses that come from a honeypot to an attacker should appear to be real. Therefore, the entire network must be simulated to ensure that the hacker will not realize that they are being tracked. According to Nathalie Weiler, there are three aspects to create a believable honeypot. The attack must be detectable, the attack can be directed to the honeypot, and the honeypot must be able to simulate the organization's known network infrastructure (Weiler, 2002). Packet filtration must take place in order to determine malicious packets from normal ones. Therefore, in order for the attack to be detectable, it must have packets that can be filtered from normal traffic. This filtration also will tie into the second aspect which ensures that the attack is directed to the honeypot. If malicious packets are identified, they should be sent to a honeypot. Lastly, the results that a hacker gets must be indistinguishable from a real network. Therefore, if any attacks are made against a system, it must appear to be giving back responses that a hacker would expect. If all of these three aspects are met, then the honeypot should work as intended, and provide further information about what the hackers are attempting to do to the system.

IP Rotation

The majority of IP's that are used are dynamic. Unlike static IP's, they change every time that the system is rebooted. Therefore, by restarting the router and changing your IP, the DDoS attack is gone. This will not make the attack go away, but you will no longer be the recipient of the malicious packets. Static IP's are very trackable. This is a major disadvantage because if one bad actor finds out your IP, they can attack it maliciously constantly. Dynamic IP's, however, change regularly on devices. This solves the aspect of traceability because nobody can find out for the long term what the systems IP is. Dynamic IP's also doesn't require individual configuration and are very easy to use.

In DDoS mitigation, this process of IP rotation could take place in varying degrees. The theory is that if there is an influx of traffic, you could preemptively change IP's to remove malicious and artificial traffic from being directed towards your systems. Given that the attack is directed towards a specific IP, there is no way that attackers could direct their attacks at a moving target. This method is especially useful if you have another IP prepared that could be utilized.

Conclusion

We believe that all organizations with internet presence should utilize DDoS mitigation. However, the ways in which DDoS mitigation may be utilized may differ based on the size of the business. While both large and small companies may use built-in protection from ISP's, the utilization of machine learning and honeypots would best be used within large businesses because of the cost associated. Smaller businesses could benefit more from rate limiting, third-party services, and IP rotation. DDoS attacks are growing in frequency and strength. Therefore, the best things that businesses can do is protect themselves from these attacks. Hospitals, banks,

and insurance companies just a few targets of these malicious attacks, and not only monetary assets are lost. If systems are not properly maintained, even lives can be lost. DDoS mitigation is the best solution possible for these attacks, and we must implement as many defenses as possible to combat these attacks.

Works Cited

Mahdi, Nikrad. “An Alternative Approach to Rate Limiting - Figma Design.” *Medium*, Figma Design, 4 May 2017, medium.com/figma-design/an-alternative-approach-to-rate-limiting-f8a06cf7c94c.

“7 DDoS Attack Trends Spotted in 2018.” *Data Foundry*, 28 Mar. 2019, www.datafoundry.com/blog/ddos-attack-trends-2018.

Atkinson, Doug. “How Machine Learning Takes Network Monitoring to the Next Level.” *Best Network Monitoring Vendors, Software, Tools and Performance Solutions*, Best Network Monitoring Vendors, Software, Tools and Performance Solutions, 31 Oct. 2018, solutionsreview.com/network-monitoring/machine-learning-network-monitoring/.

Barr, Jeff. “Reduce DDoS Risks Using Amazon Route 53 and AWS Shield | Amazon Web Services.” *Amazon*, Amazon, 20 Mar. 2018, aws.amazon.com/blogs/aws/reduce-ddos-risks-using-amazon-route-53-and-aws-shield/.

Conran, Matt, et al. “The Rise of Artificial Intelligence DDoS Attacks.” *Network World*, Network World, 11 July 2018, www.networkworld.com/article/3289108/the-rise-of-artificial-intelligence-ddos-attacks.html.

“DDoS Blackhole Routing.” *Cloudfare.com*, 2019, www.cloudflare.com/learning/ddos/glossary/ddos-blackhole-routing/.

“DDoS Mitigation | Checklist for Choosing a Mitigation Provider | Imperva.” *Learning Center*, Imperva, www.imperva.com/learn/application-security/ddos-mitigation-services/.

“Getting Started.” *Getting Started*, Rapid7, metasploit.help.rapid7.com/docs/.

"Honeypot-Based Monitoring of Amplification DDoS Attacks." *APNIC Blog*, 30 July 2018, blog.apnic.net/2018/07/30/honeypot-based-monitoring-of-amplification-ddos-attacks/.

"How AI Plays a Role in Both Stopping and Committing DDoS Attacks." *Security Intelligence*, 21 June 2018, securityintelligence.com/fight-fire-with-fire-how-ai-plays-a-role-in-both-stopping-and-committing-ddos-attacks/.

"Inside the Distil Content Protection Network." Distil.it, 2012. https://www.distilnetworks.com/wp-content/uploads/2013/09/Distil-Technical-White-Paper.pdf

Kovacs, Eduard. "Hacktivist Gets 10-Year Prison Sentence for DDoS Attack on Hospitals." *SecurityWeek*, SecurityWeek,Com, 11 Jan. 2019, www.securityweek.com/hacktivist-gets-10-year-prison-sentence-ddos-attack-hospitals.

Krupp, Johannes. "Honeypot-Based Monitoring of Amplification DDoS Attacks." *APNIC Blog*, 30 July 2018, blog.apnic.net/2018/07/30/honeypot-based-monitoring-of-amplification-ddos-attacks/.

"Metasploit: Penetration Testing Software." *Rapid7*, Metasploit, www.rapid7.com/products/metasploit/.

Netscout. "Network Security Infrastructure Report | NETSCOUT." *NETSCOUT®*, www.netscout.com/report/.

Parra-Novosad, Natalie. "7 DDoS Attack Trends Spotted in 2018." *Data Foundry*, 28 Mar. 2019, www.datafoundry.com/blog/ddos-attack-trends-2018.

"Real-Time DDoS Protection Solutions." *Corero*, www.corero.com/products/corero-smartwall-threat-defense-system.html.

Rouse, Margaret. “What Is Honeypot (Computing)? - Definition from WhatIs.com.” *SearchSecurity*, searchsecurity.techtarget.com/definition/honey-pot.

Shridhar, Kumar, and Nikhil Gautam. “A Prevention of DDos Attacks in Cloud Using Honeypot .” International Journal of Science and Research (IJSR), Nov. 2014. https://pdfs.semanticscholar.org/f49a/3c63f315c39898410f8289cdabd5cab13da3.pdf

Soenke, Justin. *Phase 3*, www.phase3.net/how-to-create-a-honeypot-to-catch-a-hacker/.

“The Advantages & Disadvantages to a Static IP Address.” *Techwalla*, www.techwalla.com/articles/the-advantages-disadvantages-to-a-static-ip-address.

Weiler, Nathalie. “Honeypots for Distributed Denial of Service Attacks.” Swiss Federal Institute of Technology ETH, 2002. http://www.csl.mtu.edu/cs6461/www/Reading/Weiler02.pdf

“What Is DDoS Mitigation?” *Cloudfare.com*, www.cloudflare.com/learning/ddos/ddos-mitigation/.