



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Implementing and Auditing the Critical Security Controls - In-Depth (Security 56
at <http://www.giac.org/registration/gccc>

Building and Maintaining a Denial of Service Defense for Businesses

GIAC (GCCC) Gold Certification

Author: Matt Freeman, matt.freeman@hawaiiantel.com

Advisor: Rob VandenBrink

Submitted: 1/14/2017

Abstract

Distributed Denial of Service (DDoS) attacks have been around for decades but still cause problems for most businesses. While easy to launch, DDoS attacks can be difficult to sustain and even more difficult to monetize for attackers. From the business perspective, a DDoS attack might result in lost revenue but is unlikely to have the same long term impact that a data breach may have. Recent changes in the IT landscape have made DDoS a more attractive attack vector for hackers. The industry trend to connect more and more devices to the Internet (often with minimal to no security), dubbed the “Internet of Things” has created a new marketplace for bad actors to sell their resource exhaustion services. Businesses need to consider all options when planning and implementing a defensive posture against denial of service attacks. As security vendors continue to offer new (and expensive) options to defend against these attacks, how does an InfoSec manager know which is best for their business. Using an “Offense informs the Defense” approach, this paper will analyze the methods used during DDoS attacks in order to determine the most appropriate defensive postures.

1. Introduction – The Need For Improved Distributed Denial of Service (DDoS) Defenses

“Time is money” is an adage that describes how important it is for a business to always be available to sell its wares. The phrase was first credited to Benjamin Franklin in the 1700s but has relevance today as businesses can maintain a 24-hour sales presence through their online operations. That is of course, if their business can stay online. The denial of service attack, a reliable old tool of hackers and cyber-criminals has seen a resurgence in popularity over the last 10 years (Anstee, 2016). Businesses are dedicating additional IT resources to plan for and mitigate denial of service attacks to ensure their business operations are always available. Time is money and unexpected downtime is one of the worst messages an IT manager will deliver to the business owners.

Estimating revenue lost to network outages caused by a denial of service attack can be difficult. Some security vendors have suggested that each hour of downtime can be worth up to \$40,000 in lost revenue (Matthews, 2014). Lost revenue will vary by business size and industry and goes beyond the loss of sales associated with an unreachable ecommerce site. Intangible costs such as damage to brand reputation can have a long-lasting impact to economic performance. Network improvements and upgrades implemented because of the attack can also add to the overall cost. The actions taken by the InfoSec manager before, during, and after a DDoS event will help to limit potential damages and protect the business from repeating the experience.

This document is presented as a reference for IT and InfoSec managers in small and medium sized businesses to understand the threats of a Distributed Denial of Service attack and how to prepare the organization to respond. For the purpose of this discussion, small/medium sized businesses (SMB) will be defined as between 50 and 500 employees with an average IT staff of less than 10 people. Annual average revenue of the business is less than \$10 million per year.

Matt Freeman, matt.freeman@hawaiiantel.com

2. Overview of DDoS Attack Methods

To prepare an effective defense against DDoS attacks it is necessary to understand how an attack works and the terminology that vendors and experts use to describe them. Using an “Offense Informs the Defense” methodology, IT professionals can analyze the common attack vectors and select the best defensive approaches. While there are many ways to describe a denial of service attack, the most common attacks are frequently referred to as volumetric or application. Each of these can be delivered in multiple ways and are often described with modifiers such as floods, reflections, or amplification.

During a distributed denial of service attack, an attacker will direct multiple devices to send data and/or make requests of a specific target device. By focusing hundreds and thousands of devices towards a single target, the attacker can amplify the impact of the attack, overwhelming the target network’s ability to respond to legitimate network requests. The objective is to degrade availability of the target.

With billions of devices connected to the Internet, DDoS attackers can build large networks of devices to use in large scale attacks. Attackers build these networks by installing applications or modifying existing code on the devices to allow them to take control. Malicious software imbedded in files attached to emails or downloaded from the Internet will install itself to a device. Once installed, the malware will attempt to connect over the Internet to a control server where it may download and install additional files. These devices remain functional and the malware attempts to evade detection by the user. An attacker that is proactively trying to build their network of devices will scan the Internet looking for IP addresses that respond to their reconnaissance tools. A commonly used application such as NMAP can report on what type of device is responding, whether the device has open ports, and what operating system is being used (Lyon, 1997). Using this information, an attacker can use any known vulnerabilities to remotely access and install their malware. These hijacked devices are referred to as bots and in large

Matt Freeman, matt.freeman@hawaiiantel.com

quantities make up a botnet. Botnets can number in the tens of thousands and wait on standby for the owner to initiate an attack on a target.

2.1. Volumetric Attacks

The volumetric attack method is one the most common DDoS attack types and one that receives the most attention (Neustar, 2016). Put simply, a volumetric attack is an attempt to overwhelm the target's network capacity to degrade availability. In the OSI model, volumetric attacks can target both layer 3 (network) and layer 4 (transport). A layer 3 attack targets the routing and switching infrastructure using common protocols such as ICMP and ARP (US Cert, 2014). Layer 4 attacks typically use TCP \ UDP protocols intending to oversaturate the Internet bandwidth of the target. Both methods can create a similar impact to the target network: the inability to access resources.

One example of a volumetric DDoS attack is an ICMP Flood. The Internet Control Message Protocol (ICMP) is a layer 3 protocol used to report errors in data processing between a packet's destination and its original source (Postel, 1981). Network administration tools such as Ping rely on ICMP to determine device availability across the Internet. The typical size of an ICMP packet should be 8 bytes, but an attacker employing an ICMP Flood attack can generate a large amount of ICMP traffic by setting a customized packet size to the maximum allowed by the target network. The attacker will then generate constant ICMP traffic from multiple sources to a single destination. If the attacker can generate enough traffic through the ICMP Flood, the target network will become unable to process legitimate network requests resulting in the loss of network availability. The ICMP Flood attack is simple to execute and available to anyone willing to spend time learning to use the Kali Linux operating system. Unsecured web-accessible cameras can be hijacked and used in ICMP Flood attacks (VandenBrink, 2016).

The capacity used in volumetric DDoS attacks has increased significantly over the last few years. In September 2016, an attack of 620Gbps sustained traffic from an estimated 24,000 Matt Freeman, matt.freeman@hawaiiantel.com

devices was directed at popular security blogger Brian Krebs (Krebs, 2016). One week later, an attack measured at 1.1Tbps of sustained traffic from 145,000 devices was launched on the web-hosting provider OVH (OVH, 2016). While these recent attacks are (currently) the high end of the spectrum, a 2016 survey of 1000 businesses reports that 63% of DDoS attacks are between 1 and 100 Gbps (Neustar, 2016).

Volumetric DDoS attacks share several key characteristics: low complexity, difficult to determine attribution, and exploitation of a common, but limited resource (network capacity). Volumetric DDoS attacks may be one of the easiest attack types to launch. An attacker needs only to know the target IP address to initiate an attack in several ways. For the no-skill attacker an Internet search on "Website Stresser" will provide dozens of DDoS 'as a service' sites. For less than \$30, a novice attacker can initiate a 3Gbps DDoS attack for 60 minutes against the target of their choice (IPStresser, 2016). For slightly more advanced attackers, toolsets such as Kali Linux or Metasploit contain multiple denial of service applications. Many of these toolsets are available for free.

Another characteristic that makes volumetric attacks popular is that the attacker can be difficult to determine. Being able to launch an attack anonymously is desirable for the attacker but requires more effort to cover their tracks. Using a website stresser service is an easy form of attack but is also easy for security analysts to trace back to the source. Criminals will often prefer to compromise legitimate servers and use as part of a larger attack (Weagle, 2016). In this example, there is no need to obfuscate where the attack is coming from since the attacker is not directly associated with the compromised system.

The third characteristic of a volumetric attack is that regardless of the attack method, the attacker is attempting to exploit the limited network capacity of the target. The Federal Communications Commission defines broadband as 25Mbps for residences (FCC, 2016). Although broadband Internet speeds are increasing in availability across the United States, the cost and time to provision are limiting factors for businesses. It is easy to see how a business's

Matt Freeman, matt.freeman@hawaiiantel.com

network can be overwhelmed when 63% of all DDoS attacks are above 1Gbps and attackers can leverage 3Gbps for as low as \$30 an hour.

2.2. Application Attacks

The application-based attacks are another method of DDoS. An application DDoS attack targets layer 7 of the OSI model to tie up resources of a specific application with the intent of preventing legitimate use of that application. Commonly attacked application protocols include HTTP, HTTPS, and DNS (Imperva, 2015). This attack method is measured in connections per second and simulates a high volume of normal user traffic to overwhelm the server that is hosting the application.

Application attacks can take many forms and the impact of this type of attack will be unique to each application. A website's search function is an example of a web-based application vulnerable to attack. For example, a single visitor to a website that sells cars may make several searches on different vehicle models and colors. Each search request might require a query of the dealership's database for vehicle availability but is within the normal operating parameters of the web server and database server. For this scenario, the database can hypothetically accommodate at maximum, 1000 requests per second. An attacker utilizing a botnet of 10,000 devices can generate similar searches forcing the search form to make 1,000s of requests per second. The search requests on the database server will exceed capacity, resulting in all user searches (legitimate and botnet) to receive null responses. Excessive use of a legitimate web service such as search forms is an example of an HTTP request Flood (Radware, 2016).

Application attacks can be difficult to detect because the attack simulates normal user activity. As web-based applications become more complex, the ability to distinguish between legitimate user behavior and malicious activity grows more challenging. Attackers using a large botnet can replicate a basic web search thousands of times per second. Conversely, an attacker

Matt Freeman, matt.freeman@hawaiiantel.com

using a tool such as Slowloris will tie up server resources from a single device. Each of these attacks is not immediately distinguishable from normal usage patterns.

Application attacks succeed through exhausting the server resources of the target. Unless the targeted system has performance monitoring in place, the administrators may not notice a steady increase in CPU or decrease in available disk space. Once the target has reached a resource limit, the application being served from it may become unavailable to normal users.

Denial of service attacks can come in many shapes and sizes but all have at least one shared objective: preventing legitimate use of the target resource. It is important for an information security manager to understand the basic attack types to prepare an effective defense. It will also be important to understand the reasons for an attack and an attacker's motivation.

2.3. Motivations for DDoS Attacks

One of the primary motivating factors for cybercrime is financial gain. Cyber-criminals employ a variety of tools to launch attacks against their victims, most with the intent of extortion or theft. Denial of service attacks are easy to launch and maintain for the attacker. For the average victim, it may be impossible to determine who was behind the attack without assistance from law enforcement agencies. DDoS attacks may not be as easy to monetize for attackers as ransomware or theft of electronic records (such as credit card information) but remains a profitable endeavor for attackers (Krebs, 2016).

Extortion is one of the most common motivators an attacker has for launching a denial of service attack. A classic example of the "protection racket" used by the American Mafia has evolved online. In the traditional definition, a criminal will approach a business with a demand for money in exchange for protecting them from other criminal activities. A similar approach is now used by DDoS attackers. Cyber gangs such as the Armada Collective and Lizard Squad will

Matt Freeman, matt.freeman@hawaiiantel.com

launch a denial of service attack against a targeted business for a short duration. This initial attack is to demonstrate to the victim their capability and preview of the impact a sustained outage will have on their business. An extortion demand is then delivered to the victim by: pay up or face extended attacks (Paine, 2016). Limited opportunity events are particularly vulnerable to this type of threat since they have a small window of opportunity to make sales. Just before the 2014 World Cup, Anonymous launched a series of attacks against World Cup websites putting millions of dollars of potential earnings at risk (Herberger, 2014).

Denial of service attacks can often be used as a diversionary tactic to mask the criminal's true intent. Often referred to as a "smokescreen", a DDoS attack against a business can divert resources away from other vulnerable areas of the network. While IT administrators are dealing with a network or server outage, attackers may have a higher chance of succeeding in installing malware or exfiltrating data from an existing compromise (Neustar, 2016).

Attackers can also have non-financial motivations. Anyone that has played an online game knows that response time from the host server is a critical component to success. Some players are willing to cheat and install third party software to help them win. Others will look to eliminate their main competition directly through out of game means, including DDoS attacks. If an attacker can determine the IP address of their competition, he or she can disable that player's ability to connect to the game. Attackers can create a competitive advantage or simply retaliate for losing to better players. As the idea of E-Sports grows more popular (and lucrative), the ability to disable an opponent before he or she can compete becomes attractive to the less scrupulous gamers (Maiberg, 2015).

Another common motivator for DDoS attacks is to retaliate against or attempt to censor opinions they disagree with. Hacktivists, a nickname for hackers that launch attacks for political and social reasons, often use DDoS attacks to disable the websites of their enemies. In a case from December 2016, anonymous attackers disabled a website belonging to the Thailand

Matt Freeman, matt.freeman@hawaiiantel.com

government. The DDoS attack was launched as a protest to what the attackers viewed as oppressive legislation limiting Internet freedoms (Ashok, 2016).

2.4. Target Selection for DDoS Attacks

Any device with an IP address that is reachable from the Internet is a potential target for a DDoS attack. A cyber-criminal launching a denial of service attack wants to make sure that the attack is highly disruptive to the target's normal operations. Attacks against business are often financially motivated with the attacker looking to create downtime for the target. Business owners will compare the lost earnings due to downtime with the financial demands of the attacker. Criminals hope that the business will pay up to avoid the lost revenue and negative publicity of an outage.

Ideal targets for DDoS attackers are highly visible, a focal point for user interaction, and critical to business operations. For most businesses, these three characteristics can describe a company's website. Websites are one of the most attacked Internet resources because they serve as the business's link to customers on the Internet (Kaspersky, 2016). Websites communicate news and information about the business. When a DDoS attack disables a website for users, the brand's reputation can be tarnished. Websites enable e-commerce for many businesses and in some cases, are the only avenue available to customers to purchase goods and services. Business websites range from simple to complex, but all rely on servers, network infrastructure and Internet RFC standards, making them vulnerable to a variety of attack types.

Criminals can also cause disruption for businesses by directly attacking their business' operations. Many small and medium businesses do not have multiple Internet gateways for their corporate network. DDoS attacks on a business network can prevent workers from reaching cloud-based applications or offsite file storages. As more businesses move into offsite-hosted solutions for email and business critical applications, an attack on the corporate Internet gateway will be very disruptive for operations. For businesses that use a VOIP-based phone system that

Matt Freeman, matt.freeman@hawaiiantel.com

relies on a hosting provider for call routing, their call center may be unable to receive and make calls to customers resulting in lost sales.

Service providers of Internet or cloud services are also ideal targets for DDoS attackers. An attack that saturates an ISP's network capacity will have a ripple effect for subscribers. In November 2015, ProtonMail experienced a sustained attack that prevented an estimated 1 million users from accessing their secure emails (Yen, 2015). Datacenters offering virtual server hosting for businesses also face a similar risk. A single vulnerability in a hosted virtual server can be exploited, impacting other subscribers on the same hypervisor.

3. Defending Against DDoS Attacks

Information Security managers can start to prepare for DDoS attacks as they would other cyber threats. The Center for Internet Security maintains a list of 20 Critical Security Controls (CSC) that can help focus defensive programs for IT security managers. The CSC use the "Offense in the Defense" approach to prioritize defensive recommendations that are based on real world attacks.

One of the most effective programs an information security team can implement is asset inventory and tracking (CIS, 2013). Critical Security Control #1 calls for the creation and maintenance of an inventory of authorized devices for use on the network. Knowing the infrastructure and critical services in a network is a pre-requisite for building defenses for them. The asset inventory should include all network infrastructure, servers, workstations, hosted applications, and business critical systems. For DDoS-specific defenses, identify systems that are most likely to be attacked such as corporate Internet gateways and websites. The IT department should regularly audit the assets to ensure no new devices were added without going through a change management process.

With an asset inventory completed, the IT team should begin building monitoring solutions for the critical systems. Critical Security Control #6 recommends the maintenance, Matt Freeman, matt.freeman@hawaiiantel.com

monitoring, and analysis of audit logs to improve anomaly detection. For networking infrastructure, such as routers and switches, use SNMP monitoring tools that can communicate with the device and report on device availability. Most network monitoring systems can be configured to alarm on predefined performance thresholds, such as high CPU usage or bandwidth usage targets. In the event of a DDoS attack, alarms from network devices indicating sustained high utilization can be an early signal to IT staff that an attack is occurring.

Servers should be monitored for availability and resource thresholds, such as CPU utilization and available disk space. Windows-based servers can utilize the SCOM system for monitoring while Linux operating systems can deploy OSSEC software. For web and file servers, IT administrators should determine normal behavior under heavy load to establish an expected performance baseline. Monitoring software can be configured to generate alarms when baseline targets are exceeded over a specified amount of time. Using the SNMP monitoring system, set threshold alarms for bandwidth utilization. If the Internet gateway circuit is provisioned with a 100 Mbps circuit, consider creating an alarm when bandwidth utilization exceeds 90% for more than 10 minutes. Watching for sustained heavy usage will help to identify a possible attack. Configure firewalls to inspect traffic being sent out (egress) along with the standard inbound inspection (ingress). If a system inside the network has been compromised and is being used as part of a DDoS attack, an egress threshold alarm on bandwidth utilization will help to identify the problem. Collecting and analyzing netflow data is also an effective method to know what applications are using the data passing through the network. Closely monitoring and alarming on network performance statistics can help detect a DDoS application attack.

With monitoring and alarming solutions in place, IT administrators can move on to implementing their defenses. One of the most effective defenses against any cyber-attack is a well-maintained IT environment. IT managers must ensure a regular patching and vulnerability assessment program is in place. Vulnerability scanning of internal and external IP ranges using automated scanning system such as Nessus or OpenVAS will help an IT administrator know which systems are vulnerable to cyber-attack. Maintaining a regularly enforced patching

Matt Freeman, matt.freeman@hawaiiantel.com

program will help the business keep a secure operating environment. While this may not specifically prevent DDoS attacks, it can help limit impact of the smokescreen attacks. As a result of using a patch management system, attackers will most likely find less systems vulnerable to attack and may not be able to establish a foothold.

In addition to vulnerability scanning of the network, IT security administrators can consider performing stress tests against probable DDoS targets. Using available tools in Kali Linux or Metasploit the administrators can simulate a DDoS attack against themselves to determine how their systems perform while under attack. Another benefit of stress testing the network and webservers is to help determine at what level of attack the network begins to fail. Knowing the limits of defenses may help to justify future upgrade expenses if defenses are not able to maintain service under moderate attack levels. As with all intrusive testing, ensure executive approval has been received and documented. Perform intrusive testing outside of business operating hours to limit impact to the business.

A controlled stress test against a company's own network will help to prepare the IT administrators for an actual attack. Using the data collected in the test, staff can validate their alarming thresholds for critical devices. Systems that were impacted without generating alarms can be tuned to help with detection. Attack symptoms should be documented and transitioned into an Incident Response Plan for the organization. An IR plan will act as a playbook for staff to follow during and after an attack. The IR plan should include documentation on the procedures to follow during an attack, roles and responsibilities, and a communication plan. The communication plan should include all necessary IT staff as well as executives, public relations, and operations managers. Communication plans should also include external contact information and expectations. External communication contacts may include an emergency escalation point at the Internet service provider or information for law enforcement agencies. Who needs to be contacted and when should be determined before an attack occurs so that IT staff can react correctly during an incident.

Matt Freeman, matt.freeman@hawaiiantel.com

Creating an asset inventory then implementing and testing defenses is part of training the IT staff. The IT manager should designate a resource to convert all the preparation materials into a play book that contains the procedures to be followed during a DDoS attack. Use the play book to improve team readiness by performing reviews of the procedure with employees. As devices are added to and removed from the network update the play book. Be sure to review and update contact information every few months.

During an active denial of service attack, a prepared business should be able to quickly detect the event and react per their Incident Response Plan. IT staff should know the procedures for mitigating the attack based on the play book training. Mitigation options during an active attack depend on the network infrastructure in place and the agility of the business to implement alternatives. One of the most important steps during an attack is to determine what type of attack is happening and then to identify the source of the attack. IT administrators should review firewall and server access logs to understand the attack patterns and likely targets. Where possible, block the originating IP addresses using firewall access control lists. Attacks using large botnets will easily overwhelm an IT administrator's ability to find and block IP addresses but the information gathered during the log analysis will help. If the DDoS attack is aimed at a corporate Internet gateway, switching to backup circuits through another ISP may be able to restore some connectivity to the network. Networks without a redundant Internet connection will need to contact their ISP to request assistance. Provide the ISP's support staff with the source IP information and request they implement routing rules that will prevent the traffic from reaching the target network.

If a DDoS attack succeeds in disabling web services for an extended amount of time, the business should consider their backup options. Building redundancy in critical services is an action to perform before an attack occurs. For example, if a website is under sustained attack and the business has prepared an alternate hosting site with identical content, IT administrators can change the "A" record of the website and redirect users to the secondary servers. While the attacker can also adjust to the new target, the website can now be accessed at multiple hosts

Matt Freeman, matt.freeman@hawaiiantel.com

which will force the attacker to split their attack between them. The web security company Cloudflare offers DDoS mitigation solutions for websites that can be implemented in just over 30 minutes (Zatlyn, 2016).

While IT administrators are addressing the impact of the DDoS attack, IT managers can begin following the communications plan. Keeping internal users aware of the situation does not require a step-by-step recap of the troubleshooting and mitigation steps but should instruct department managers on how their staff can continue working. Public relations staff can begin drafting notifications to news sources to provide reasons for the unplanned downtime.

Once the attacks have subsided, the IT staff should immediately begin searching for indicators of compromise across the network. DDoS attacks are often used as a diversionary tactic and are combined with other cyber-attacks (Kaspersky, 2016). If the network has monitoring already in place across the network, recent alarms should be reviewed for anomalous activity. Mirroring network traffic to a network IDS from a core network switch can provide monitoring of internal network traffic. Watching internal network traffic is an effective detective measure that can assist with identifying unauthorized network activity. In addition to reviewing IDS alarms, IT staff should review server and firewall logs for signs of compromise. If the business has a SIEM collecting logs from all IT assets, then correlation of events becomes much more efficient.

Cyber-attacks against a business are sure to prompt questions from the owners. Information security managers should expect to provide a summary of the event to executives. Managers should strive to complete analysis of the event within a few days. Using the information collected, managers should create an after-action report that includes details and metrics about the event in terms the business owners can understand. The report should begin with a summary of the incident that explains what happened. Post-incident reports should contain the start and end time of the event and a list of impacted services. Alarms generated by internal monitoring systems will document how long it took the IT staff to detect the incident. Notes

Matt Freeman, matt.freeman@hawaiiantel.com

should be used from the incident to determine when the response began. IT managers can use the time to respond metric as a baseline for future incidents and work with the team to improve the response time. This report should focus on technical and procedural facts. The report should include which systems experienced degraded availability and for how long. After the report has been presented and discussed with senior management, review the findings with the IT support teams. Analyze the key metrics and discuss areas that the incident response and play books can be improved. If the network defenses were insufficient to withstand the DDoS attack use the report's data to highlight areas that need improvement.

3.1. DDoS Mitigation Options

As the frequency and strength of DDoS attacks increase, the number of vendors offering defensive solutions has also increased. Each vendor solution offers a different way of identifying and mitigating attacks to limit impact to an organization. The right solution for a network depends on a variety of factors, such as the number of business locations, the amount of cloud-hosted applications, your business's tolerance for downtime, availability of high speed Internet services, and existing network infrastructure. For most small and medium businesses, the deciding factor often comes to down IT budget. Without making vendor recommendations, this paper will explore what services are protected best by which solution and how well they mitigate volumetric and application attacks. This paper will examine three categories of DDoS mitigation: on-premises solutions, upstream / transport solutions, and DDoS mitigation “as a service” solutions.

On-premises solutions for DDoS mitigation are designed around devices that are physically installed on the network. This category includes equipment such as UTM firewalls, web application firewalls, and DDoS mitigation appliances. On premises solutions provide the most control over mitigation defenses for the IT administrators because they have direct administration of the device. On-premises solutions can be overwhelmed by high capacity

Matt Freeman, matt.freeman@hawaiiantel.com

volumetric attacks. For example, a network administrator can configure a rule on the UTM firewall to deny all traffic from a specific IP address during an attack. Unfortunately, during a DDoS attack, botnets can leverage thousands of IP addresses from all around the world. Even if administrators can keep up with creating deny rules, the firewall itself is still processing the traffic before rejection. The Internet gateway will be overwhelmed by incoming attack traffic and legitimate network traffic will be impacted.

Web application firewalls can be deployed on the local network between the Internet gateway and the web servers. If properly configured, this firewall can be effective at mitigating application attacks against the webserver but will likely still fail during volumetric attacks. These devices require extensive baselining and tuning to be able to detect application attacks. Administrators need to understand normal web server activity over time so they can implement rules that prevent or alarm on unusual traffic.

DDoS mitigation appliances work by examining patterns and protocol ratios within network traffic to determine whether network activity is an attack or not. Vendors have created complex algorithms that can identify what a normal amount of protocol traffic would be for a specific amount of data. For example, 1 GB of data can be expected to have a certain amount of UDP, TCP, and NTP traffic with an expected size per protocol per packet. During an ICMP Flood attack the DDoS mitigation appliance should detect that an unusual amount of ICMP traffic is being processed and that the traffic is larger than normal. The appliance will drop abnormal traffic and allow legitimate traffic to be processed.

Another category of DDoS defenses is implemented upstream from the target network and attempts to filter or redirect attack traffic so that it never reaches the target. This service is frequently sold by Internet service providers as an enhanced offering to their standard Internet transport. Upstream mitigations can be applied by the ISP upon customer request which can cause a delay in restoring normal traffic. The ISP support staff may require their customer (who is under attack) to specify what IPs and protocols to filter, putting an additional delay on

Matt Freeman, matt.freeman@hawaiiantel.com

mitigation. A DDoS attack against a specific subscriber may impact all subscribers if the ISP does not have enough bandwidth capacity to manage the attack. In some cases, the ISP may blackhole traffic for the attacked user to protect other subscribers. Blackholing traffic sends all traffic for that customer to a non-routable destination, effectively shutting down the Internet service. Upstream filtering services can be a low cost and scalable solution for volumetric attacks but offers limited assistance during application attacks. It is necessary for the business to know what solutions their ISP provides and at what cost. The IT Manager can weigh the costs of upstream filtering against the potential loss due to DDoS disruption and act accordingly.

Another defensive option that works specifically for websites and is both an upstream and “as a service” solution. Referred to as Content Delivery Networks, websites are replicated across servers in geographically diverse areas. The primary objective of a CDN is to improve the loading speed of a webpage being accessed by a visitor to the website. This happens because Internet users typically have a lower latency for websites hosted closer to their location. Business websites built on a CDN have the additional benefit of some immunity to volumetric DDoS attacks since the website is not being served from a single server location and is more resilient to DDoS. One limitation is that websites supported by a CDN often need to communicate with a single database server to return search results or authenticate user login requests. If a DDoS attacker can identify a single point of failure on a CDN-supported website some functions of the site may become unavailable. In December 2016, Amazon Web Services (AWS) announced AWS Shield, a DDoS defense for subscribers using the AWS hosting platform (Barr, 2016). As the largest cloud hosting provider in the world, AWS provides a robust option for businesses looking for mitigation solutions.

Any IT manager that has attended a vendor conference has heard the term “as a service” in a dozen different contexts. Denial of service mitigation “as a service” describes an off-site hosted mitigation solution with a subscription fee attached. DDoS mitigation services operate similarly to on-premises solutions but in an upstream location. Network traffic being routed to the target business network is inspected and assessed for potential attack patterns. During an

Matt Freeman, matt.freeman@hawaiiantel.com

attack, network traffic is “scrubbed” so that attacking traffic is dropped while legitimate traffic is passed to the business network (Radware, 2016). DDoS Mitigation as a Service solutions can protect against both volumetric and application attacks.

3.2. Business Justification for DDoS Defenses

Businesses face a variety of cyber threats for which they must prepare defenses. Malware and ransomware threats require an antivirus end-point solution. Malicious attachments and unsolicited email require mail filtering solutions. Protecting the corporate network against intrusion requires many layers of defense such as firewalls, IPS, IDS, and SIEM solutions. Defending the network requires multiple layers against a seemingly endless number of attack vectors. How does the IT manager account for the expense of building and maintaining a DDoS defense when there are so many other systems to plan for?

IT managers should complete a risk assessment to determine the likelihood of their network being a target. In addition, they should work with business owners to estimate potential losses and costs associated with downtime from DDoS attack. Solutions provider, Imperva, has created a model that helps businesses estimate their risk of attack based on their industry and attack surface (Incapsula, 2014). For example, a traditional brick-and-mortar only business is less likely to be attacked than an online gaming company. Another factor that increases risk of a successful attack is a network's surface area which is measured by the number of devices reachable from the Internet. IT managers need to know which of their critical business systems can be impacted by a DDoS attack and how losing access to those systems will impact business operations. Using the estimates for lost revenue and associated costs due to DDoS attack, the IT manager can have an informed discussion with senior management about how a DDoS mitigation solution fits into the company's annual budget.

4. Conclusion

Matt Freeman, matt.freeman@hawaiiantel.com

Prepare for increase in DDoS attacks

The frequency and strength of DDoS attacks are expected to continue to increase in the future. In 2010, the largest recorded DDoS attack was measured at 100Gbps (Anstee, 2016). November 2016 saw an attack over 1Tbps in size. Bandwidth availability to businesses is increasing and costs are coming down, but few businesses are prepared to defend against an attack of 1Tbps. Manufacturers continue to produce low cost devices with Internet connectivity built in that are being exploited by hackers (Krebs, 2016). So far in 2017, there are few signs that indicate vendors are willing to increase the security of their products. Improving device security would increase production costs and consumers tend to prefer inexpensive gadgets.

Information security managers have a variety of defensive tools available to them. Knowing which tool works best in a given environment depends on how well the IT manager knows their business operations. Select the right defensive solutions to limit surface area vulnerable to attack. IT Managers will reduce the overall time to respond to an event by providing the IT staff with training and monitoring tools ahead of time.

As attackers continue to find ways to monetize DDoS activity, the need for businesses to build and maintain an effective DDoS defense increases. Whether the attacker is using DDoS as an extortion threat or just a distraction, the potential impact to business remains high. In most cases, IT managers that have completed a risk assessment of their vulnerability to DDoS attack will find that the cost to implement a defense is less than the potential long term costs of a sustained outage.

Matt Freeman, matt.freeman@hawaiiantel.com

References

- Anstee, D., Bowen, P., Chui, C.F., Sockrider, G. (2016). Worldwide Infrastructure Security Report Volume XI. Retrieved from https://www.arbornetworks.com/images/documents/WISR2016_EN_Web.pdf
- Matthews, T., (2014, November). DDoS Impact Survey Reveals the Actual Cost of DDoS Attacks. Retrieved from <https://www.incapsula.com/blog/ddos-impact-cost-of-ddos-attack.html>
- Lyon, G., (1997) NMAP Reference Guide. Retrieved from <https://nmap.org/book/man.html>
- Neustar (2016, October). Worldwide DDoS Attacks & Protection Report. Retrieved from <https://www.neustar.biz/resources/whitepapers/ddos-attacks-protection-report-us-2015>
- US Cert (2014). DDoS Quick Guide. Retrieved from <https://www.us-cert.gov/security-publications/DDoS-Quick-Guide>
- Postel, J. (1981, September). RFC 792: Internet Control Message Protocol. Retrieved from <https://tools.ietf.org/html/rfc792>
- VandenBrink, R. (2016, July). Pentesters (and Attackers) Love Internet Connected Security Cameras!. Retrieved from <https://isc.sans.edu/forums/diary/Pentesters+and+Attackers+Love+Internet+Connected+Security+Cameras/21231/>
- Krebs, B. (2016, September). KrebsOnSecurity Hit With Record DDoS. Retrieved from <https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/>
- OVH. (2016, October). The DDoS that didn't break the camel's VAC*. Retrieved from <https://www.ovh.com/us/news/articles/a2367.the-ddos-that-didnt-break-the-camels-vac>
- IPStresser (2016). Purchase. Retrieved from <https://www.ipstresser.com/index.php?page=purchase>
- Weagle, S (2016, June). Cyber Criminals Sell Compromised Servers to Carry Out DDoS Attacks. Retrieved from <https://www.corero.com/blog/734-cyber-criminals-sell-compromised-servers-to-carry-out-ddos-attacks-.html>
- FCC. (2016, January). 2016 Broadband Progress Report. Retrieved from <https://www.fcc.gov/reports-research/reports/broadband-progress-reports/2016-broadband-progress-report>

Matt Freeman, matt.freeman@hawaiiantel.com

Imperva. (2015). The Top 10 DDoS Attack Trends. Retrieved from https://www.imperva.com/docs/DS_Incapsula_The_Top_10_DDoS_Attack_Trends_ebook.pdf

Radware. (2016). DDoS Attack Definitions – DDoSPedia. Retrieved from <https://security.radware.com/ddos-knowledge-center/ddospedia/http-flood/>

Krebs, B. (2016, September). Israeli Online Attack Service ‘vDOS’ Earned \$600,000 in Two Years. Retrieved from <http://krebsonsecurity.com/2016/09/israeli-online-attack-service-vdos-earned-600000-in-two-years/>

Paine, J. (2016, April). Lizard Squad Ransom Threats: New Name, Same Faux Armada Collective M.O. Retrieved from <https://blog.cloudflare.com/lizard-squad-ransom-threats-new-name-same-faux-armada-collective-m-o-2/>

Herberger, C. (2014, June). Anonymous Delivers DDoS to the 2014 World Cup. Retrieved from <https://blog.radware.com/security/2014/06/anonymous-delivers-ddos-to-the-2014-world-cup/>

Ashok, I. (2016, December). Hackers hit Thai government with DDoS attacks protesting against restrictive internet law. Retrieved from <http://www.ibtimes.co.uk/hackers-hit-thai-government-ddos-attacks-protesting-against-restrictive-internet-law-1597339>

Maiberg, E. (2015, August). eSports Has a DDoS Problem. Retrieved from <http://motherboard.vice.com/read/esports-has-a-ddos-problem>

Kaspersky. (2016, February). Businesses Don’t Need a Website to Be a Victim of DDoS – Attacks Target Internal Systems Too. Retrieved from http://newsroom.kaspersky.eu/en/texts/detail/article/kaspersky-lab-businesses-dont-need-a-website-to-be-a-victim-of-ddos-attacks-target-internal/?no_cache=1&cHash=5c3786c08d36fd2d19615f6094b6a8ca

Yen, A. (2015, November). DDOS Updated. Retrieved from <https://protonmaildotcom.wordpress.com/2015/11/05/ddos-update/>

Center for Internet Security. (2013, September). Top 5 CIS Controls. Retrieved from <https://www.cisecurity.org/critical-controls.cfm>

Zatlyn, M. (2016, December). I am under DDoS attack, what do I do? Retrieved from <https://support.cloudflare.com/hc/en-us/articles/200170196-I-am-under-DDoS-attack-what-do-I-do->

Kaspersky. (2016, May). Research reveals hacker tactics: cybercriminals use DDoS as smokescreen for other attacks on business. Retrieved from <http://newsroom.kaspersky.eu/en/texts/detail/article/research-reveals-hacker-tactics->

Matt Freeman, matt.freeman@hawaiiantel.com

[cybercriminals-use-ddos-as-smokescreen-for-other-attacks-on-business/?no_cache=1&cHash=28d60d3b5e4523f0b3956850de5cec3b](https://www.sans.org/whitepapers/cybercriminals-use-ddos-as-smokescreen-for-other-attacks-on-business/?no_cache=1&cHash=28d60d3b5e4523f0b3956850de5cec3b)

Barr, J. (2016, December). AWS Shield – Protect your Applications from DDoS Attacks. Retrieved from <https://aws.amazon.com/blogs/aws/aws-shield-protect-your-applications-from-ddos-attacks/>

Radware. (2016). DDoS Attack Definitions – DDoSPedia. Retrieved from <https://security.radware.com/ddos-knowledge-center/ddospedia/scrubbing-center/>

Incapsula. (2014) DDoS Downtime Cost Calculator. Retrieved from <https://lp.incapsula.com/ddos-downtime-cost-calculator.html>

Anstee, D., Bowen, P., Chui, C.F., Sockrider, G. (2016). Worldwide Infrastructure Security Report Volume XI. Retrieved from https://www.arbornetworks.com/images/documents/WISR2016_EN_Web.pdf

Krebs, B. (2016, October). Hacked Cameras, DVRs Powered Today's Massive Internet Outage. Retrieved from <https://krebsonsecurity.com/2016/10/hacked-cameras-dvrs-powered-todays-massive-internet-outage/>

Matt Freeman, matt.freeman@hawaiiantel.com

Upcoming Training

Click Here to
{Get CERTIFIED!}



Community SANS Ottawa SEC566	Ottawa, ON	Sep 25, 2017 - Sep 29, 2017	Community SANS
Rocky Mountain Fall 2017	Denver, CO	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Phoenix-Mesa 2017	Mesa, AZ	Oct 09, 2017 - Oct 14, 2017	Live Event
SANS Tysons Corner Fall 2017	McLean, VA	Oct 14, 2017 - Oct 21, 2017	Live Event
SANS Tokyo Autumn 2017	Tokyo, Japan	Oct 16, 2017 - Oct 28, 2017	Live Event
SANS San Diego 2017	San Diego, CA	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Gulf Region 2017	Dubai, United Arab Emirates	Nov 04, 2017 - Nov 16, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DC	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Cyber Defense Initiative 2017 - SEC566: Implementing and Auditing the Critical Security Controls - In-Depth	Washington, DC	Dec 14, 2017 - Dec 18, 2017	vLive
SANS Security East 2018	New Orleans, LA	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZ	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS New York City Winter 2018	New York, NY	Feb 26, 2018 - Mar 03, 2018	Live Event
Community SANS Ottawa SEC566	Ottawa, ON	Feb 26, 2018 - Mar 02, 2018	Community SANS
San Francisco Spring 2018 - SEC566: Implementing and Auditing the Critical Security Controls - In-Depth	San Francisco, CA	Mar 12, 2018 - Mar 16, 2018	vLive
SANS San Francisco Spring 2018	San Francisco, CA	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS 2018	Orlando, FL	Apr 03, 2018 - Apr 10, 2018	Live Event
SANS OnDemand	Online	Anytime	Self Paced
SANS SelfStudy	Books & MP3s Only	Anytime	Self Paced