State of the Art of DDoS Mitigation Techniques

Franz Josef Ennemoser, Patrick Sattler* and Johannes Zirngibl*
*Chair of Network Architectures and Services, Department of Informatics
Technical University of Munich, Germany
Email: franzjosef.ennemoser@tum.de, sattler@net.in.tum.de, zirngibl@net.in.tum.de

***Abstract*—Distributed Denial of Service (DDoS) attacks continue to be one of the biggest threats for online services. This has created a large demand for DDoS Protection Services (DPS) in the last decade, who use their clouds to defend customers from larger attacks every year. Since these types of attacks are launched from many different sources, preventing or mitigating DDoS attacks requires sophisticated defence mechanisms. The paper shows the three basic components of a potent defence against DDoS attacks which are attack detection, traffic classification and attack response. While the defence mechanisms of DPS providers are proprietary, we showcase some mechanisms that demonstrate how DPS systems can be comprised in practice. Furthermore, we explore the current leading vendors of DDoS protection systems such as Akamai, which is responsible for serving 15 to 30 percent of the world wide web traffic, or the well-known company Cloudflare, that offers unmetered DDoS protection even for their free plans.**

***Index Terms*—Denial-of-service, distributed denial-of-service, distributed denial-of-service mitigation, DDOS, networks**

1. Introduction

The Internet is becoming increasingly important to society as billions of devices are now networked and more are being added every day. This increasing importance of accessibility means that there is also an ever greater incentive to disrupt it. One of the most common attacks on online services are Distributed Denial of Service (DDoS) attacks, which have become more frequent and more intense in the last two decades. Large attacks in Q1 2020 were again breaking records in peak bandwidth with Amazon reporting a 2.3 Terabits/s attack on their AWS servers [1]. Comparing this to the 5-6 Terabit/s average bandwidth of Frankfurt's internet exchange point DE-CIX in Q1 2020 reveals the size of such an attack [2].

A DDoS attack tries to overload a service with various methods with the goal of the service being unable to answer requests of legitimate users. A DDoS attack is a special kind of DoS (Denial of Service) attack, in which the source of the attack is distributed over multiple devices that cooperate to overwhelm the targeted service. The attack devices are often botnets, which are networks of compromised computers under the control of the attackers. Creating such a network is often accomplished by infiltrating computers through the usage of malware such as trojans and worms.

This stealing of computing and network resources already creates a great imbalance in expenses between attacking and defending side. That is an invitation for many attackers to attack their targeted web services in order to inflict financial damage as well as harming the public image. As defending against such sophisticated attacks that grow in size every year is no easy task, this has given a rise in popularity of DDoS protection provider to hide web services behind or in their large cloud networks [3].

In Section 2 the taxonomy of DDoS attacks is explained with basic examples. The three components of a DDoS defence system: attack detection, traffic classification and attack response are discussed in Section 3. Section 4 compares current leaders in the market of DDoS protection services and outlines the increasing adoption of such services. In Section 5 the paper is concluded and future work is mentioned.

2. Types of Attacks

DDoS attacks can be grouped into different categories and these diverse types of attacks call for different defence mechanisms. Furthermore, there are attacks called multi-vector attacks which try to combat this by combining several attack techniques. An unfortunate property of an online service is the fact that successfully attacking the weakest link in the network can stop the whole network. A strong DNS server will not help in the case the webserver itself is overloaded with dummy requests from the attacker [4]. While not all attacks can be categorized perfectly, there exist three basic types of attacks:

2.1. Volume-based Attack

This is the most common type of attack. The targeted network node is attacked with a sheer amount of dummy requests created by the botnet controlled by the attacker with the goal of depleting the available network bandwidth. This results in legitimate traffic being unable to pass through and the service is taken down.

Since volume-based attacks need large bitrates to be successful, a common attack is the DNS amplification attack. This attack abuses the fact, that DNS requests can receive large answers compared to the size of the request. This amplification can reach an amplification factor of 50+. Since DNS is UDP-based, the attacker uses the victims IP address as source address, which in turn will be targeted by a large amount of DNS packets as can be seen in figure 1.

doi: 10.2313/NET-2022-07-1_01

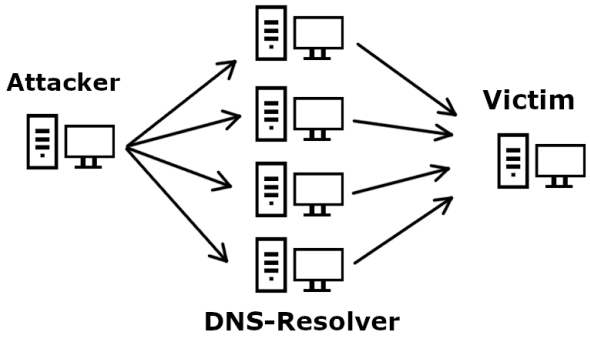

Figure 1: Congesting the network uplink of the victim with a DNS-based flood attack. (Reworked from Loukas et al. [5])

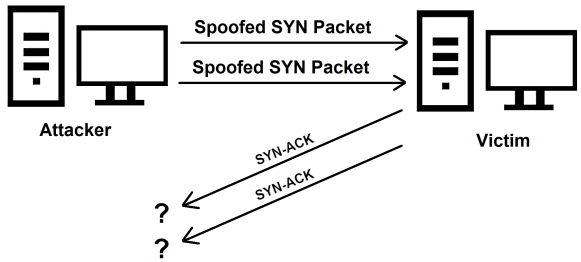

Figure 2: Depleting the resources of the victim with a TCP-SYN attack. (Reworked from [6])

2.2. Protocol-based Attack

These attacks mainly use vulnerabilities and shortcomings of protocols in OSI layers 3 and 4 to exhaust processing or memory resources of the target node instead of the network bandwidth. To be able to measure and compare the strength of these attacks, packets per seconds (pps) are usually used as metric.

A prominent example of this is a TCP SYN attack which exploits the way a TCP connection is established. The attackers send requests with active SYN flags to which the server responds with an SYN-ACK. Usually, the client would acknowledge this but in this case the attacker does not answer at all as it can be seen in Figure 2. Now the server has to wait for the ACK to timeout, wasting valuable memory space. To increase the difficulty of defending against this attack, the attacker will most likely also spoof the source IP address of the SYN packets.

2.3. Application Layer Attack

Attacks in this category are destructive compared to the small effort on the attacker side. They are also harder to detect as they often closely resemble real user behaviour.

An example for an application layer attack is the Slowloris attack. It abuses the HTTP protocol by sending incomplete HTTP GET requests without termination code and refreshes the connection just before the server would timeout the corresponding session. Over time this will occupy all connections the server is able to open at the same time and consequently service is unreachable for all users.

3. DDoS Mitigation Techniques

DDoS Mitigation can happen at various locations in the network. While we want to be as close to the source as possible to prevent the malicious network load from reaching large parts of the network, the distributed aspect of the attack makes this a challenging task. In a real-world scenario a service operator has the choice between three basic operating approaches. He can either run his own mitigation solution, outsource it to a DDoS Protection Service (DPS) provider or use a hybrid approach combining both solutions, each with its own benefits and downsides.

For most attacks the DPS should include the following three components [5]:

- **Detection:** First step in mitigating an attack is the simple detection of said attack. Attacks became more sophisticated and therefore distinguishing flash events from DDoS attacks has become harder. Detection can be grouped into anomaly-based and signature-based detection systems. While in anomaly-based detection the DPS has to first learn the normal user behaviour and later on detect an abnormal deviation to that, the signature-based detection tries to fit current observations into known patterns to detect attacks.
- **Classification:** As soon as an attack has been detected, the next step is to classify the incoming traffic into legitimate and invalid (created by the attacker) traffic.
- **Response:** After the malicious traffic has been marked the DPS needs to drop the invalid packets, preferable at the network edge to be less affected by the massively increased network traffic.

3.1. Detecting a DDoS Attack

The detection of the attack is the first step to be able to act on it. While it may sound like a simple comparison between normal traffic and the high volume traffic of an attack, there are also legitimate events that generate a high-volume of traffic. This could be an announcement, a product release or a news article linking to the specific service. In that case, dropping packets could heavily impact the companies behind the webservice either financially or in public reception. Furthermore, the attacks themselves are evolving and emerge in different shapes and sizes. Just by polling and comparing the traffic bandwidth alone it will be difficult to recognize an application-layer attack that does not rely on a huge attack bandwidth.

A common way to distinguish detection methods is to classify them either as anomaly-based or signature-based [3]:

Signature-based detection compares current network traffic with known attack behaviour, resulting in a high detection ratio and low false-positive rate. This is only true for known patterns and will be relatively ineffective for newly emerging attacks.

Anomaly-based detection will in most cases have a higher false positive rate, but also be effective in detecting new types of attacks. Anomaly-based detection divides

doi: 10.2313/NET-2022-07-1_01

further into statistical analysis and machine learning based detection. In statistical analysis the system observes metrics such as packet arrival rate, packet type arrival rate and entropy of packet header fields. While it provides fast detection rates, without further inspection the false positive rate may be high. In learning-based systems, data mining techniques highlight previously unknown connections in the incoming traffic [3].

3.2. Classifying DDoS Traffic

A closely connected step is the classification of the incoming traffic, to be able to separate any legitimate traffic from the DDoS attack packets. As with detection, the classification works in a signature- or anomaly-based fashion and compares features to usual known traffic patterns. Features can be real-time gathered statistical features or even actively created by letting the user prove their legitimacy. This task is usually done by *dedicated validity tests*, which can be passive or active [5].

3.2.1. Passive Validity Tests.

Loyal clients: A very basic thought example is that a user that requests a news site every day, can be regarded as highly trustworthy even during times of an attack. This works even in cases when attackers spoof the source IP address of their packets, as it is unlikely, that they will randomly find a trusted IP address.

Time-to-live: Even though source addresses during an attack can and will often be invalid, the hop count of an IP packet will give insight into how far the packet travelled. This is done by Hop-Count filtering introduced by Jin et al. through comparing the TTL of the incoming packet with a table that stores known mappings between IP-addresses and their hop-counts [7]. It might not be a working traceback method but it can be compared to the apparent source IP address and give hints about the legitimacy of the packet.

3.2.2. Active Validity Tests.

Active validity tests in comparison are in direct communication with the user of the incoming request and challenge it to prove its legitimacy. This has become an important step since sophisticated attacks started to imitate a natural increase in bandwidth similar to external events. Therefore, classifying this kind of traffic has become harder. Active tests work on the premise that legitimately increased traffic patterns will be created by humans instead of automated programs. Famously used for this task are Reverse Turing tests, such as CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart).

While CAPTCHAs come in different types, all of them are based on challenges trivially solvable by a human but hard to solve by a computer. Challenges range from tasks like reading obscured digits and letters to classifying a group of images. In 2013 Google also included a behaviour-based analysis of the browser interactions as a filter in its CAPTCHA service called reCAPTCHA. This improves usability for deemed low-risk users as they are not tasked to solve time-consuming challenges. In case the fingerprinting does not rate the user as credible, another verification with classic challenges is added successively.

3.3. Responding to an DDoS Attack

The DDOS attack responses are typically classified depending on their location in the network. They are classified in source-based, network-based and destination-based techniques each with its own advantages and disadvantages as shown by Dietzel et al. [8]. Due to the distributed aspect of DDoS attacks, the easiest position to detect an attack is directly at the target (destination-based) but the mitigation is less effective. Mitigating an attack close the source would be ideal but is difficult to realise in practice as attacks can be launched from anywhere. The following paragraphs outline a selection of techniques used in practice:

A destination-based mitigation technique is dropping the packets that have been marked high-risk by the classification component. As mentioned in the last section, these filters work time-based, history-based or hop-count-based.

Especially with volume-based attacks, filtering alone will not alleviate the pressure on the network resources. Even if the heuristics would allow for a good classification, the number of incoming packets would still overload the target network. For this reason and as an additional mitigation step, mechanisms such as adaptive rate limiting by Ioannidis et al. [9] or IP traceback by Adler et al. [10] are proposed.

Adaptive Rate Limiting is based on the concept of aggregates which are subsets of the traffic that share a common property. These properties include the packet destination, the type of packet and packets with a bad checksum. If an aggregate responsible for a significant portion of the traffic is found, the aggregate is propagated to upstream routers to rate-limit the malicious traffic. Traffic adhering to the rate-limit will be allowed while other traffic will be dropped to as mentioned by Zargar et al. [11].

IP Traceback mechanisms try to find the true sources of the forged IP packets. Since IP routing is stateless and routers usually only know where to forward the incoming packet, the routers have to support the traceback method to be able to contribute to IP traceback mechanisms. The main categories of traceback techniques are packet marking and link testing. In packet marking the routers add their identification to the packet probabilistically in order to enable the victim to identify the path of malicious traffic after receiving enough packets. In link testing the routers closest to the victim get iteratively tested until the source of the attacker's traffic can be reached [11]. The effectiveness of IP traceback in practice is limited, since the traceback mechanism would need to be deployed with minimal cost in time and storage, low false positive rate and while respecting privacy of the inspected packets [5]. Attackers can also forge their own marked packets and therefore disturb the traceback mechanism.

Increasing the Attack Surface may sound a little contra-productive as reducing it is an important step in many cyber security related topics. The nature of DDoS attacks however concentrate on a single point which makes mitigating the attack almost impossible if the attack surface is reduced to a certain point [3]. Therefore increasing

doi: 10.2313/NET-2022-07-1_01

the attack surface is a common strategy to help mitigate an DDoS attack or at least soften the impact on the targeted network. In practice this is usually achieved by using cloud providers that hide the user either in or behind their large networks.

An effective last resort response to a DDoS attack is the so-called *blackholing* of the traffic flowing to the victim. Blackholing is defined as dropping the traffic at the routing level, which can be implemented at almost any router with no additional performance impact. The victim autonomous system announces the prefix to blackhole to its upstream network via BGP (Border Gateway Protocol) which nowadays means the blackholing will atleast happen at one of the supporting IXPs (Internet Exchange Point) [8]. Blackholing is usually the last resort since the announced prefixes will by unreachable by both the attackers and the legitimate users, but it will reduce collateral damage to neighbouring devices and networks.

4. DDoS Protection Service Provider

Many web services today rent a cloud service to allow them to cost-effectively scale their operation. Many of the cloud service provider also offer their cloud as a DDoS Protection Service. The traffic will be sent through the cloud where the traffic will be cleaned from malicious packets in so called scrubbing centers. The clean traffic will be rerouted to the webserver responding to the requests of valid users.

DDoS protection via a cloud service can either be always-on, on-demand or a hybrid version that combines both. The most common always-on solution is the usage of a CDN (Content Delivery Network), which distributes the content over many cache servers around the world to be geographically close to the end user. CDNs are not only used for QoS (Quality of Service) objectives, they can also be used as DDoS protection as the distribution of content reduces the effects of an attack [3].

In case an on-demand strategy is desired, a reactive plan that reroutes the traffic only in the case of an attack to the cloud is appropriate. The cloud then scrubs the incoming packets and sends back the clean traffic to the webserver. The routing of the traffic can either be done by making a change in the DNS record of the victim or by a BGP advertisement change.

4.1. DPS Provider Overview

In Forrester Research' 2021 report of DDoS Protection Service provider 11 significant vendors are mentioned and compared against each other [12]. This paper focuses on the 4 leaders in the market according to Forrester Research which are Akamai, Cloudflare, Imperva and Radware. While the vendors keep their filter technology and scrubbing techniques proprietary, the mechanisms are not significantly different compare to on-premise detection [4].

Akamai is one of the largest cloud service providers, with a network of over 300,000 servers in 135 countries serving between 15% and 30% of the web traffic. This large network size accumulates to a network capacity of more than 175Tbps. Akamai is targeted towards enterprise customers, as it has a minimum contract of 12 months and does not reveal pricing information without requesting a quote [12].

Cloudflare is another big CDN provider with a strong focus on DDoS mitigation. In comparison to the enterprise focused Akamai, Cloudflare offers start with a free plan containing the option of adding additional features via their Pro and Enterprise plans priced $20 and $200 per domain respectively. Their basic volumetric DDoS protection is already included in the free plan, which is also unmetered. Defence mechanisms against layer 7 attacks have to be purchased additionally as a package [13].

Radware is an Israelian company that offers application delivery and several cybersecurity products. In contrast to Akamai and Cloudflare, Radware is one of the oldest and largest vendor for on-premise DPS devices, but they have been transitioning towards cloud-based and hybrid approaches in the last years. Due to their long presence in the industry they have a deep understanding of DDoS attacks. Therefore, they are especially suited for difficult attack cases [12].

Imperva is another security specialists that offers cloud-based DDoS protection services. Imperva advertised it's large network size in the last years but according to [12], most of their competitors have caught up and even surpassed the capability of Imperva's network. The capability to deflect even the largest attacks can currently still be found at all large cloud providers.

4.2. DPS Adoption

The increasingly large DDoS attacks every year also increase the pressure on web service provider to employ a cloud-based DPS to be able to mitigate them. Jonker et al. have proposed a methodology to check domain names for active traffic diversion to a cloud-based DPS and used it to analyse all .com, .net and .org TLDs containing over 50% of the names in the global namespace with daily snapshots over 1.5 years between March 2015 and September 2016. While the amount of domains in that namespace grew 9% from 140M to 152M domains, the number of domains protected by the top 9 leading DPS provider grew 24% to a total of about 9M domains. For 6 months they have also been monitoring the Alexa Top 1M list as well as the .nl TLD, whose DPS usage grew 12% and 11% during that time respectively [14]. This shows clearly the increased interest in the services of DPS providers.

5. Conclusion and Future Work

In this paper we provided an overview of basic DDoS attack types and general mitigation strategies. This was followed by an overview of leading vendors in the DPS market. Thereafter, the ongoing trend towards cloud-only or hybrid-based DDoS Protection Services was outlined. They are becoming the most popular remaining option to have a chance against these Terabits per second large volumetric attacks, which have occurred increasingly often in the last few months and years. Future challenges and opportunities lie in the field of SDN (software-defined networking) and the usage of machine learning based detecting and filtering of malicious traffic.

doi: 10.2313/NET-2022-07-1_01

References

[1] "AWS Shield - Threat Landscape Report," May 2020. [Online]. Available: https://aws.amazon.com/de/blogs/security/aws-shield-threat-landscape-report-now-available/

[2] "DE-CIX Traffic Statistics." [Online]. Available: https://www.de-cix.net/en/locations/frankfurt/statistics

[3] I. Ozcelik and R. Brooks, *Distributed Denial of Service Attacks: Real-world Detection and Mitigation*, 05 2020.

[4] E. Chou, R. Groves, and a. O. M. C. Safari, *Distributed Denial of Service (DDoS): Practical Detection and Defense*. O'Reilly Media, 2018. [Online]. Available: https://books.google.at/books?id=G19PwAEACAAJ

[5] G. Loukas and G. Öke, "Protection Against Denial of Service Attacks," *Comput. J.*, vol. 53, no. 7, p. 1020–1037, Sep. 2010. [Online]. Available: https://doi.org/10.1093/comjnl/bxp078

[6] "SYN-Flood-Attack, Cloudflare Learning." [Online]. Available: https://www.cloudflare.com/de-de/learning/ddos/syn-flood-ddos-attack/

[7] C. Jin, H. Wang, and K. G. Shin, "Hop-Count Filtering: An Effective Defense against Spoofed DDoS Traffic," ser. CCS '03. New York, NY, USA: Association for Computing Machinery, 2003, p. 30–41. [Online]. Available: https://doi.org/10.1145/948109.948116

[8] C. Dietzel, A. Feldmann, and T. King, "Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild," in *International Conference on Passive and Active Network Measurement*. Springer, 2016, pp. 319–332.

[9] J. Ioannidis and S. Bellovin, "Implementing Pushback: Router-Based Defense Against DDoS Attacks," 03 2002.

[10] M. Adler, "Tradeoffs in Probabilistic Packet Marking for IP Traceback," *Journal of the ACM (JACM)*, vol. 52, no. 2, pp. 217–244, 2005.

[11] S. T. Zargar, J. Joshi, and D. Tipper, "A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks," *IEEE communications surveys & tutorials*, vol. 15, no. 4, pp. 2046–2069, 2013.

[12] D. Holmes, "The Forrester Wave: DDoS Mitigation Solutions, Q1 2021," *Forrester Research*, March 2021.

[13] "Cloudflare Pricing and Plans." [Online]. Available: https://www.cloudflare.com/plans/#overview

[14] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras, "Measuring the Adoption of DDoS Protection Services," ser. IMC '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 279–285. [Online]. Available: https://doi.org/10.1145/2987443.2987487