Exploring DDoS Defense Mechanisms

Patrick Holl
Betreuer: Oliver Gasser
Seminar Future Internet SS2014
Lehrstuhl Netzarchitekturen und Netzdienste
Fakultaet fuer Informatik, Technische Universitaet Muenchen
Email: holl@in.tum.de

ABSTRACT

Nowadays, Distributed Denial-of-Service (DDoS) attacks are a major threat for all sizes of networks. The number of attacks against companies and institutions steadily increased over the last years. Downtime of an enterprise network usually causes financial damage. Therefore, it is important to have mechanism for DDoS defense. In this paper, various DDoS defense mechanisms are reviewed and compared with focus on rule and model based approaches. Large Botnets allow for new kinds of attacks like flash crowd simulation which mimic a huge mass of organic traffic. These kind of attacks are difficult to detect and new defense techniques are required. In order to discover new mitigation algorithms, it is necessary to understand at which layers attacks can happen. Therefore, we take a look on how attacks are classified in current research literature. In addition to the attack classification, rule and model based DDoS defense mechanisms are reviewed. For both model and rule based techniques scenarios exist where one algorithm outperforms the other one. Having this in mind, we list the advantages and drawbacks of both techniques based on insights of research literature. Emerging architectures like SDN may change the way DDoS defense is handled. Researchers are already working on algorithms that are suitable in SDN environments. The goal of this paper is to summarize current defense mechanisms and give a brief outlook on how DDoS defense could look like in the future.

Keywords

DDoS attacks, DDoS Defense, DDoS Mitigation, Algorithm comparison

1. INTRODUCTION

Denial-of-Service (DoS) attacks are a major threat for the availability of the global internet infrastructure. The main goal is to limit or even prevent intended users to access a service. In most cases, an attacker controls several compromised machines which are distributed over the internet. Such distributed attacks are also called Distributed Denial-of-Service (DDoS) attacks. Lately, attack networks with over 400,000 compromised machines were revealed [1]. A huge Botnet like this is able to cause severe availability problems even to large web services. Depending on the offered service, downtime can cause loss of revenue or other negative effects for the one who runs the service. As a consequence, defense mechanisms to detect and mitigate such DDoS attacks are necessary. DDoS defense is an active field of research but also the attackers evolve their tools and algorithms to overcome detection and mitigation. In this paper, we want to give an overview of several DDoS attacks on the one side and defense algorithms on the other side. Professional DDoS attacks often aggregate traffic from their compromised machines in a way that it looks like organic traffic from intended users. Attacks that mimic natural users can be particularly difficult to detect. Lately, new ideas emerged how to do DDoS defense in modern network architectures like Software-defined networking (SDN). However, no studies about how the proposed mechanisms work in real world environments exist nowadays.

The increasing complexity of DDoS attacks requires many-faceted defense mechanisms. Therefore, modern defense systems make use of several detection and mitigation techniques. DoS attacks can be handled in various ways, e.g. by building an infrastructure around the service which is able to survive a DDoS attack by deploying resources dynamically based on the packet load the service gets. Reactive defense mechanisms are on the other hand algorithms that try to detect and mitigate attacks at the time they occur. The reactive approaches are classified as rule or model based. Model based approaches check for traffic anomalies and rule based ones for certain patterns, e.g. in a specific packet header field.

Not all defense mechanisms are suitable for all kinds of DDoS attacks. In some cases, rule based algorithms can outperform statistical approaches, for instance, when the setup time must be very short. However, there are also scenarios where model based algorithms have advantages over rule based ones, e.g., in blocking Zero-Day DDoS attacks. Zero-Day DDoS attacks are not yet publicly known attacks.

In section 2 of this paper, DoS and DDoS attacks are defined in more detail. Section 3 gives an overview of DDoS defense mechanisms that emerged over the last years in research literature. In section 4 we compare rule based and statistical approaches and state the advantages and shortcomings of each approach. The last section 5 gives an outlook on the future of DDoS defense in SDN. SDN is an emerging technique in networking but has not yet replaced traditional architectures. DDoS defense in SDN is an active area of research right now but no studies are available yet that prove or falsify the concepts and hypotheses of the researchers.

doi: 10.2313/NET-2015-03-1_04

2. DEFINING DOS AND DDOS

To be able to mitigate DoS and DDoS attacks, respectively, it is necessary to understand the difference between both attacks. In the following section, the differences and characteristics of those kind of attacks are described. Furthermore, the basic structure of a DDoS attack is analyzed. In the last part of this section, DDoS attacks are classified based on the information of current research.

2.1 DoS vs. DDoS

The primary goal of DoS attacks is to make a service or the whole network unavailable to its intended users. To this end the DoS attack targets a network node to hinder it from processing packets that originate from legitimate requests [2].
A Distributed-Denial-of-Service (DDoS) attack can be seen as a special form of a DoS attack. In this case, distributed means the usage of multiple machines to attack the target [3]. This basic difference is visualized in Figure 1.

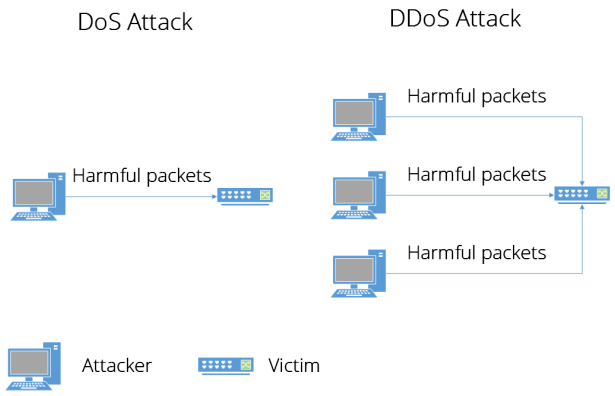

Figure 1: DoS vs. DDoS

2.2 Basic structure of a DDoS attack

Figure 1 shows only a very simplified view of a DDoS attack. The actual structure of a typical attack is more complex. Attacks that are able to take down large web services typically need several thousand compromised machines. For example, 2008 a Botnet consisting of over 400,000 machines was revealed [1]. The coordination of such a huge, distributed attack network is complex and typically done in a three layered structure as described by Kelm et al. in [4].

On the first layer is the attacker itself who controls several handlers on layer two. The handlers on layer two are used to automatically compromise and control machines to act as agents on layer three. To compromise the agents, the handlers use automated routines to find and exploit vulnerabilities. One handler can control hundreds of agents which are then used to send harmful packets to a victim. In Figure 2 we can see how the traffic, separated in control and attack parts, flows within the DDoS attack structure.

2.3 DDoS classification

To be able to develop and understand defense mechanisms, it is necessary to understand the different tiers on which a

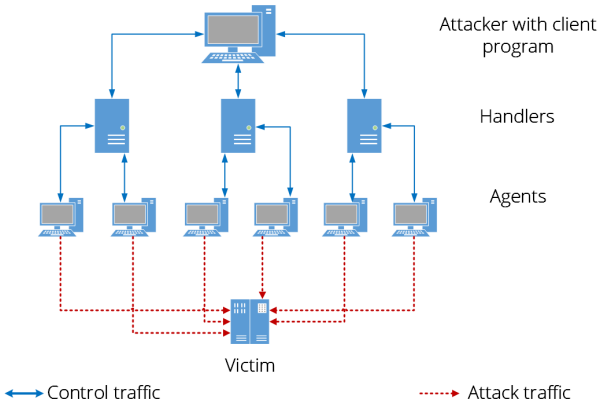

Figure 2: DDoS Control Layers

DDoS attack can happen. Several possible classifications exist nowadays which are described in detail in [5] by Douligeris et al. In the following section, we focus on the generalization of possible attack types stated in [6]:

Bandwidth-based (flood) attacks
The attacker uses its agents to send a mass amount of junk IP traffic to the victim. Consider a scenario where a website running on an arbitrary HTTP server is attacked. The webserver can only handle a certain amount of users or to be more specific, HTTP requests. If the attacker can saturate that maximum number of requests, the webserver is not able to respond to legitimate requests anymore.

Transport Layer Attacks
Protocol attacks exploit vulnerabilities and features in certain protocols. One of the most common DoS protocol attacks is TCP SYN flooding [7]. The attack exploits a weakness of the three-way handshake which is necessary to set up a TCP connection between two hosts. A valid three-way handshake consists of three messages which are sent between the client and the server. The last message is normally sent by the client and acknowledges the connection with the server. An attacker now drops the last message, meanwhile the server is waiting for the response and is – depending on its implementation – blocked for new TCP connections. Many more protocol attacks exist nowadays as described in [8] and [5].

Application Layer Attacks
Another kind of DoS attacks are targeting the application layer. An example for such a DoS attack is shown in [9] by Kulkarni where the target is the popular Apache2 webserver. Application layer attacks can target any application that is reachable via a network, e.g. expensive database requests. The underlying protocol of the applications is secondary but most commonly the HTTP protocol is used.

Actual attacks are often not easily classifiable because they exploit characteristics of more than one type. Application layer and protocol attacks often comes hand in hand with bandwidth-based attacks.

doi: 10.2313/NET-2015-03-1_04

3. DDOS DEFENSE MECHANISMS

Denial-of-Service attacks can cause severe damage on the infrastructure of the attacked victim [10]. Consider an e-commerce company which sells products online. Any downtime of the website means loss of revenue, since no legit users are able to use the service. Therefore, it is necessary to develop systems and algorithms to mitigate DDoS attacks and their impact on a service.

In the following section, we take a look at various DDoS mitigation techniques and the technical challenges that comes hand in hand.

3.1 Defense approaches

In [11], Zhang et al. categorizes DDoS defense and detection into three basic categories.

Proactive defense mechanisms

In 2002, Keromytis et al. [12] proposed a method which actually does not tackle a DDoS attack directly but built the infrastructure in a way that it will survive a DDoS attack. This implies that the attacked victim needs access to resources that can handle and survive a DDoS attack. Nowadays, such infrastructures could be called Cloud or Cloud-hosting where resources are only extended when needed. Such an infrastructure can be the only method to survive so called Zero-Day DDoS attacks, which are attacks that are not yet publicly known – and therefore no defense mechanism is available.

Reactive defense mechanisms

The concept behind reactive defense mechanisms is to mitigate or block a DDoS attack when it happens. This can be a challenging endeavor since the attack must be observable by certain patterns. An Intrusion Detection System (IDS) works as a traffic monitor and analyzer [11]. Thus means, that the DDoS defense is only as strong as the deployed IDS. Nowadays, for many DDoS attacks exist mitigation techniques for example TCP SYN Flooding [13] or ICMP Flooding [14].

Post attack analysis

The main goal of post attack analysis is to analyze an attack and find patterns in it to feed the IDS with, and on the other hand, to trace back the attacker [11]. Song et al. presented in [15] a method to trace back a spoofed IP address to its real source. However, Zhang et al. showed that it is not feasible to trace back large Botnets at the moment of the attack. One reason is that large, modern Botnets consists of thousands of agents and second one is that the global internet is too big that all administrators can collaborate to exchange trace back information.

All of the named defense approaches can be combined and applied together. For example, selective blackholing (see section 3.2) as a reactive approach to mitigate the attack itself and a dynamic *cloud* that can supply additional computation power on demand. The reactive mechanism helps to mitigate the attack in such a way that less additional resources are necessary to handle and survive the attack.

3.2 Selective Blackholing

A classical blackholing approach can be used to block traffic which is destined to a certain victim [16]. Therefore, all packets from a certain IP address that causes high traffic are routed to a so called null route. Depending on the geographical region of the source address, the traffic could theoretically also originate from legitimate request, e.g. due to some advertisement. Having this in mind, one major drawback of this approach is that not only malicious but also legitimate traffic is filtered out. In order to tackle these shortcomings, selective blackholing emerged.

Selective DDoS blackholing is a two-step process with the goal of sending all DoS related packets to a static route defined on the edge network routers to drop them [17]. In the first step of the process, all edge routers are initialized with a so called *blackhole* destination. All packets forwarded to this destination are separated from regular traffic and usually dropped. In the second step, the BGP routers in the network use the specified blackhole destination to forward packets and instruct the service provider when certain conditions are met [18]. A possible condition would be for example, a malformed packet or an IP address which is known for being an agent in a large Botnet.

In 2014, Snijders presented a selective blackholing approach which also takes the geographical scope into account [18]. Consider the following scenario: A web shop that only sells and ships products to German addresses. Most likely the customers of this shop will access it with a *German* IP given by their ISP. A large Botnet is usually distributed over several countries because the agents are (in most cases) infected by automated routines that exploit vulnerabilities in the system. A selective blackhole that takes the geographical scope of German IP addresses into account can now be used to block traffic outside this scope. Figure 3 illustrates this case.

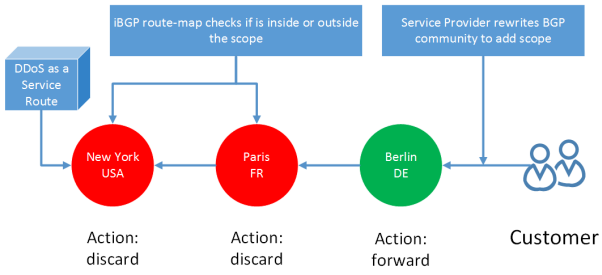

Figure 3: Selective blackholing, discard outside Germany

doi: 10.2313/NET-2015-03-1_04

Router name	Continent ID	Country code	Metro ID	Latitude, Longitude
r1.tky.jp	3	392	46	35.65671„139.80342
sjo.us	1	840	29	37.44569,-122.16111
dal.us	1	840	30	32.80096„-96.81962

Table 1: Router geographical locations table [18]

Snijders used the following four rules to illustrate his algorithm [18]:

* discard traffic sourced outside 'this' country (5580:664)
* discard traffic sourced outside 'this' continent (5580:660)
* discard traffic sourced outside a 1000 km radius from 'here' (5580:663)
* discard traffic sourced outside a 2500 km radius from 'here' (5580:662)

According to Snijders, the *this* and *here* keywords are points that refer to the point where a customer interconnects with the service [18]. The two-numbered code in brackets stands for a autonomous system and an action. For instance, *5880:664* means that the AS with the code 5580 wants to discard all traffic outside the country where a customer interconnects with the service. This action is represented by the second code 664. A router can only set packets on a null route if they have a route map where they can check whether the destination is outside the geographical scope or not. Therefore, a table which contains the geographic location of the routers is necessary. Table 1 shows how such a database could look like. In this table, column one states the name of the routers, column two to four are geographic indicators. For instance, the *Metro ID* with the value 46 represents a number code for Tokyo, Japan. The last column contains geographic coordinates for distance calculation.

Packets with source addresses that are routed through routers outside the geographical location defined by the service, can now be set on a null route and discarded.

However, selective blackholing as described above is not able to block a DDoS attack completely. When the attack traffic originates from an IP address which is within the defined scope, the traffic would still reach its target. But since only traffic from its main target group reaches the service, the attack can be heavily mitigated and the service can continue its business. Anyhow, scope based selective blackholing also has some shortcomings which we will discuss in section 4 in more detail.

3.3 Statistical Approaches

Statistical approaches are based on the assumption that DDoS attack traffic shows anomalies in the entropy and frequency of selected packet attributes. In 2003, Feinstein et al. proposed an algorithm to detect DDoS attacks by measuring statistical properties in packet headers at different points in the network [19].

A mandatory basis of every statistical detector is the model on which it is built on. For instance, the model can be generated based on a certain number of legitimate requests within a defined time range. Assume a web service provider that logs all consecutive packets from 9.00PM to 9.15PM for one month. After that month the service provider is able to build a model (for the given time range) that contains information about the distribution of the source IP addresses. In the second month, incoming packets can be checked against the model and classified as forward or drop.

3.3.1 Entropy of consecutive packets

One method proposed by Feinstein et al. is based on the entropy comparison of consecutive packet samples to identify changes in their randomness [19]. Information entropy is the average amount of information in each sample and defined as follows, where H is the entropy, n the number of symbols and p_i the occurrence probability of symbol i:

$$H = -\sum_{i=1}^{n} p_i \log_2 p_i \tag{1}$$

The source IP address is one field in the packet header that can be used to identify deviations in the randomness in comparison with legitimate requests. Depending on the number of agents, the attacker only has a limited number of IP addresses he can use. Therefore, attacks can be identified if the number of unique IP addresses has a wide variance from the legit samples.

One shortcoming of this technique is that an attacker who knows how the algorithm works is able to break it by slowly forging packets until they match the right entropy levels. However, this is not a trivial task since multiple detectors can be chained together which makes it harder to break through all of them.

Detection Quality
Feinstein et al. evaluated their proposed algorithm by simulating a DDoS attack based on an excerpt of 1,000,000 packets from the NZIX dataset[19]. They decomposed the packets into 75% legitimate and 25% DDoS traffic. TCP SYN flooding is used as an attack. Therefore, the 25% attack traffic consists of TCP SYN flooding packets. The packets used for the attack were numbered from 700,000 to 800,000. Figure 4 shows the result that the researchers got from the attack by applying an entropy model using the source IP addresses as stated above. The calculated entropy values are mapped to the y-axis and the packet count, in thousand, is mapped to the x-axis.

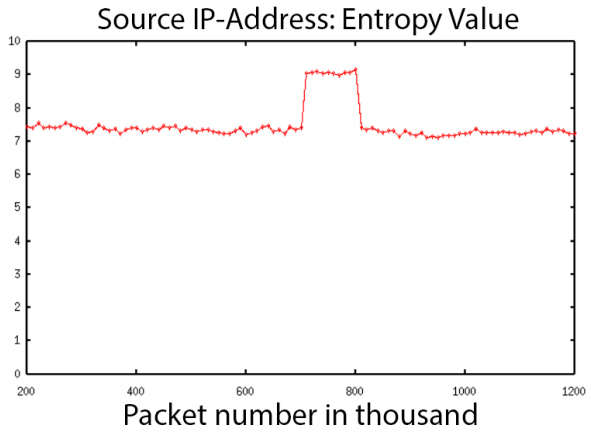

Figure 4: DDoS entropy result (Source: [19])

3.3.2 *Chi-Square Statistic*
If the number of measurement values is small, e.g., a binary value which is either 1 or 0, the entropy might not be sufficient enough to calculate reliable thresholds. As a consequence, Feinstein et al. also made use of the Chi-Square statistic. Therefore, Feinstein looked at the TCP SYN flag distribution of consecutive incoming packets. In some scenarios, the source address distribution is not an appropriate base. For instance, when Network Address Translation is used to map several source addresses onto one unique address. The TCP SYN flag is a discrete value, it is either 1 for set or 0 for unset. Pearson's chi-square Test is a suitable method to compare the distribution of discrete measurement values [19].

Pearson's chi-squared test is defined as follows, where B is the number of cells (e.g. 2 for the TCP SYN flag values), N_i the number of packets where the corresponding values occur and n_i is the expected number of packets under a normal distribution.

$$\chi^2 = \sum_{i=1}^{B} \frac{(N_i - n_i)^2}{n_i} \tag{2}$$

Detection Quality
Feinstein et al. used the same setup for the chi-square test as for the entropy test [19]. Regarding the DDoS detection quality, both techniques offer the same accuracy (see Figures 4 and 5). The only difference is that the thresholds are different to classify a packet as a DDoS packet. In this case, all packets with a χ^2 value over around 1,500 can be considered as harmful.

3.3.3 *Conclusion*
Model based statistical approaches have one major advantage over rule based techniques like selective blackholing. They allow detection for Zero-Day DDoS attacks. However, it can be a challenging task to set up an appropriate model in practice. For instance, a DDoS detector for a website with a constantly growing user base would classify legitimate requests as attack if the underlying model is static. Therefore, the model has to be constantly updated which can be tricky since the update must happen when the server is not under attack. Otherwise the model is falsified and not suitable for attack detection anymore. Another issue of model based techniques is to find the right thresholds that classifies anomalies. If the thresholds are set too low, many intended users are blocked but if the thresholds are too high, the DDoS attack can cause more damage.

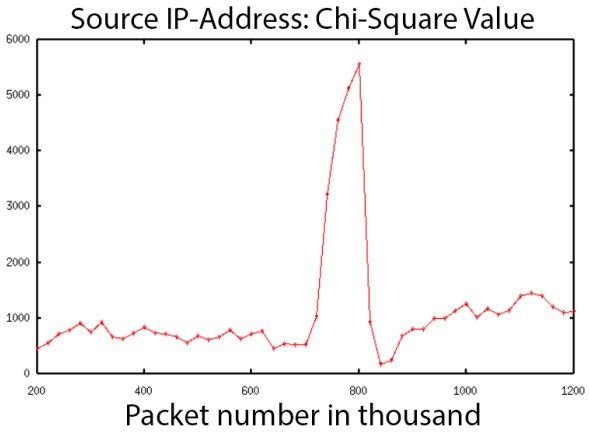

Figure 5: DDoS entropy result (Source: [19])

3.4 Challenges
As we have seen in section 3.3.3, it is challenging to provide a model based DDoS detection for an evolving website. In fact, many more technical challenges exist to detect and mitigate actual DDoS attacks. In the following section, we will discuss what DDoS attacks makes so hard to mitigate nowadays.

3.4.1 *Size of the Botnet*
Attackers who control large Botnets with tens of thousands of compromised machines can go beyond flooding attacks like TCP SYN flooding. A modern DDoS technique is to mimic a flash crowd. In this case, a flash crowd is a huge number of intended users that is accessing a web service due to some suddenly appeared popularity. For example, a blog article which is linked on the front page of a major

news site can cause high traffic for the blog. The detection of such an attack is not trivial since both legitimate and malicious requests do not differ in their content but only in their intention.

3.4.2 Abnormal traffic detection

Despite the fact that we can use statistical approaches as stated in Section 3.3 to detect anomalies in a consecutive packet flow. It is still a very challenging task to identify packets forged by an attacker which make use of several obfuscation techniques. Flash crowd imitation is one example for a technique which is hard to defend. Detection on the one hand is relatively easy since the number of requests just increases up to an abnormal state. But defense on the other hand is hard because the attacked service provider might block a real flash crowd and permanently disgruntle intended users with it. These kind of attack is very depending on the resources an attacker can use for his attack. A Botnet with a high geographical distribution is, for instance, difficult to mitigate with selective blackholing since the attacker can use bots that are within the geographical fence.

3.4.3 Long-term attacks

On the long-term, an attacker may harm a service more if he uses attacks that not completely prevents intended users from using it but increases its respond times. The goal of these attacks is to utilize a service to its capacity. They are both, hard to detect and hard to mitigate, because at the first appearance a service looks like in its default state. The only difference is that due to its high utilization a service has a longer respond time. Services that are time critical can take severe damage from such attacks. For instance, Amazon.com[1], which is one of the largest e-commerce provider on the planet, estimated a loss of 1,600,000 USD in sales if its page load time would increase by one second for the period of one year [20].

3.4.4 Large-scale testing

It is a game of cat-and-mouse between the ones who develop attacks and the ones who develop defense algorithms. The development of DDoS defense mechanisms is a complicated task because it is hard to test the developed algorithms in real-world scenarios. One reason for that is the lack of large-scale testbeds, another one is that it is not safe to perform experiments within the actual internet infrastructure [8]. Nowadays, common methods are small-scale test setups and simulations as we have seen in section 3.3.

On a commercial level, several service providers exist which offer large-scale tests. Not all of them are reliable, e.g. if they don't do any site owner verification. Such providers fall in the category of DDoS-as-Service as described in [21]. Reliable companies that offer large-scale tests don't do this for free usually. Instead they charge a price that relates with the size of the attack. As a consequence, large-scale testing can be very expensive.

[1] http://amazon.com

4. COMPARISON OF DDOS DEFENSE TECHNIQUES

In the last section we have seen several DDoS defense techniques like rule based or model based filtering. Depending on the concrete attack, one algorithm can outperform the other one. The different defense techniques can also be combined in order to increase the DDoS mitigation level. In this section we want to discuss and compare the advantages and disadvantages of the proposed techniques in section 3.

4.0.5 Advantages of rule based filtering

As described in [18], rule based filtering can be an effective way to mitigate certain DDoS attacks. In comparison with statistical approaches, rule based filters don't require a model to detect attacks. Building accurate distribution models for parameters like the source addresses can be a difficult and time consuming task. Rule based filters on the other side require much less setup time. They can start working immediately after they are setup i.e. as soon as the rules are made. Furthermore, rule based filtering allows a detection rate of 100% if a certain attack happens for which rules are already defined. In addition to that, the number of false-positive results is (depending on the rules) very low. The maintenance of a rule based filter usually requires less effort than a model based one. One reason for this is that rule based filtering is independent of the number of packets and the traffic. In comparison with that, statistical models must be constantly updated in order to fit the parameter distribution of the legitimate requests.

4.0.6 Disadvantages of rule based filtering

Application layer attacks as described in section 2 usually exploit vulnerabilities or software design mechanisms. In order to block such attacks, the filter rules must match the attack pattern. However, for unknown vulnerabilities no such rules can be defined which means that Zero-Day DDoS attacks cannot be blocked by rule based filtering. The only possibility to mitigate Zero-Day DDoS attacks is to set up generic rules like selective blackholing [18]. In this case, selective blackholing can be seen as a generic rule because packets are not further analyzed but blocked based on their geographic origin only. Nowadays, many different types of DDoS attacks exist which target the victim on different layers, see section 2. Usually, different DDoS attacks require different detection rules which results in a large repository of rules that is required to block those different kind of attacks. In comparison with that, model based approaches identify harmful packets on their abnormal parameter distribution - without having different rules for any single attack.

4.0.7 Advantages of model based filtering

One major advantage of model based approaches is that Zero-Day DDoS attacks can be mitigated. This is because traffic streams which have a high deviation with respect to the model built from the legitimate requests are flagged as potentially harmful. In addition to that, one model can be used to mitigate different kinds of attacks. For instance, an arbitrary attack where the attacker sends consecutive packets which have an abnormal source address distribution. In this case, it is not necessary to further analyze the payload of the packets themselves.

4.0.8 *Disadvantages of model based filtering*

Statistical approaches suffer from the so called cold-start problem. It takes time to build an appropriate model before attacks can be detected and mitigated. In addition to that, the models have to be constantly updated in order prevent a high rate of false-positives when the website is evolving (i.e. getting more traffic by intended users) over time. As a consequence, it can take time until the models reflect the actual situation. However, most websites are not evolving in a speed which invalidates a model very fast. Another difficulty in defining model based filters is finding the right thresholds. On the one hand, if the threshold is set too low, legitimate requests are blocked and therefore intended users, on the other hand, if the threshold is set too high, malicious requests are not blocked and therefore harm the service.

4.0.9 *Conclusion*

Both rule based and statistical approaches have advantages over the other technique. Depending on the concrete scenario, one technique can outperform the other one. Services that require a more robust detection for Zero-Day DDoS attacks should use a model based mitigation technique. Services that require a very short setup time should primarily apply a rule based filter. Since there exists a plethora of different kinds of DDoS attacks, one detection algorithm alone may be not sufficient enough. Therefore, it is possible to chain and combine several defense techniques like selective blackholing and a source address distribution model with concrete thresholds. Consider the following scenario:

A German online shop which has 8600 customers in total and 8500 of them live in Germany. A majority of 95% of all orders is shipped to a German address.

In this scenario, Germany is the main market and responsible for most parts of the revenue. A selective backholing algorithm that blocks all packets from source addresses outside Germany could be the first line of defense. Attacks from a globally distributed Botnet are severely mitigated in this case. For attacks launched by German hosts, the source address distribution of consecutive packets can be used as a second line of defense and to block abnormal packet streams. All in all we can say that a plethora of different kinds of attacks and threats like Zero-Day DDoS attacks require the combination of various defense mechanisms to take advantage of their specific strengths.

5. FUTURE TRENDS

In this section we will discuss future trends in networking and how they possibly affect DDoS defense and attack mechanisms. Vykopal et al. represents the hypothesis that SDN is ideal for distributed DDoS detection and mitigation [22].

The traditional network architecture which is based on TCP/IP is now over 20 years old but still the major technique for transmitting packets in a network. Due to trends like Cloud and the Internet of Things, the demands on network technology is constantly increasing. In this case, Cloud stands for centralized, outsourced service providers that offer disk space and applications as a service. One technique which is emerging over the last years is SDN or Software-defined networking. Figure 6 shows a schema of a SDN architecture with its three layers. SDN allows Cloud providers to easily separate the traffic from their customers into flows.

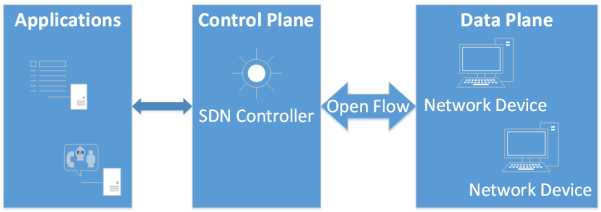

Figure 6: Software Defined Networking architecture schema

One major point that brings Vykopal to his hypothesis is that a SDN is flow based. The data plane and the control plane are separated from each other. A switch in a traditional network contains both control and data planes whereas switches in a SDN only contain the data plane. The packets are forwarded based on the entries of the flow table managed by the SDN controller. Since the traffic is organized in flows through the SDN, there is a possibility that attack flows can be identified by certain patterns. Another point of Vykopal which makes SDN ideal for him concerning DDoS mitigation is that there is a central point of knowledge, i.e. the SDN controller. Once a malicious flow is identified within the network, the controller can block or blackhole the respective flow.

At the time of this paper no studies that prove or falsify the hypothesis of Vykopal et al. were available. However the researchers will focus on three main research questions and try to answer them over the next three years [22]. The first one is a generic investigation of the differences that SDN brings to traditional networks and its monitoring. In a second step, the researchers try to explore the specific vulnerabilities in the data and control plan of a SDN. Furthermore, Vykopal et al. wants to use that discovered knowledge afterwards to find out how DDoS attacks in Software Defined Networks can be optimally mitigated.

6. CONCLUSION

DDoS attacks are one of the largest threats for the global internet nowadays. The attacks can be used to slow or even shut down large network infrastructures. Therefore, DDoS defense is a necessary task to ensure the availability of the internet. In this paper, we tried to give an overview of various kinds of DDoS attacks and how they can be detected and mitigated. Steadily, new kinds of attacks on the one hand and new defense mechanisms on the other hand are discovered. Having this in mind, it is mandatory to constantly update attack patterns and signatures for detection and mitigation purpose. Another very important point is to develop algorithms which are able to mitigate Zero-Day DDoS attacks, so that at least the main intended user group is still able to access the service. In 2014, Snijders presented a technique called selective blackholing which we discussed in section 3 [18]. Depending on the concrete scenario, selective blackholing can be a very strong defense against DDoS

attacks. However, it still has some drawbacks, e.g. if the users are not locally concentrated. As we have seen in section 4, selective blackholing could be combined with a model based algorithm in order to harden its defense abilities. In the future, multiple lines of defense can play a much more important role. This is mainly because DDoS attacks are evolving by getting more complex and resources behind it. However, SDN can dramatically change the way DDoS defense is done. For now, we don't have any major studies on this and it is an ongoing field of research as we have seen in section 5. As a consequence, the superiority of SDN in DDoS defense remains speculation at the time of this paper. In conclusion, further research in SDN DDoS defense is necessary.

7. REFERENCES

[1] "Spam on rise after brief reprieve." `http://news.bbc.co.uk/2/hi/technology/7749835.stm`. Accessed: 2014-08-30.

[2] E. Y. Chen, "Detecting dos attacks on sip systems," in *VoIP Management and Security, 2006. 1st IEEE Workshop on*, pp. 53–58, IEEE, 2006.

[3] A. Asosheh and N. Ramezani, "A comprehensive taxonomy of ddos attacks and defense mechanism applying in a smart classification," *WSEAS Transactions on Computers*, vol. 7, no. 7, pp. 281–290, 2008.

[4] K. Möller and S. Kelm, "Distributed denial-of-service angriffe (ddos)," *Datenschutz und Datensicherheit*, vol. 24, no. 5, pp. 292–293, 2000.

[5] C. Douligeris and A. Mitrokotsa, "Ddos attacks and defense mechanisms: classification and state-of-the-art," *Computer Networks*, vol. 44, no. 5, pp. 643–666, 2004.

[6] NSFOCUS, "Introduction to ddos attack," 2004.

[7] W. M. Eddy, "Tcp syn flooding attacks and common mitigations," 2007.

[8] J. Mirkovic and P. Reiher, "A taxonomy of ddos attack and ddos defense mechanisms," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 39–53, 2004.

[9] P. Kulkarni, *Responsive System for DDoS Attack against Apache Web Server*. PhD thesis, NATIONAL INSTITUTE OF TECHNOLOGY KARNATAKA, 2010.

[10] M. Sachdeva, G. Singh, K. Kumar, and K. Singh, "Measuring impact of ddos attacks on web services," 2010.

[11] G. Zhang and M. Parashar, "Cooperative defence against ddos attacks," *Journal of Research and Practice in Information Technology*, vol. 38, no. 1, pp. 69–84, 2006.

[12] A. D. Keromytis, V. Misra, and D. Rubenstein, "Using overlays to improve network security," in *ITCom 2002: The Convergence of Information Technologies and Communications*, pp. 245–254, International Society for Optics and Photonics, 2002.

[13] H. Wang, D. Zhang, and K. G. Shin, "Detecting syn flooding attacks," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3, pp. 1530–1539, IEEE, 2002.

[14] L. Limwiwatkul and A. Rungsawang, "Distributed denial of service detection using tcp/ip header and traffic measurement analysis," in *Communications and Information Technology, 2004. ISCIT 2004. IEEE International Symposium on*, vol. 1, pp. 605–610, IEEE, 2004.

[15] D. X. Song and A. Perrig, "Advanced and authenticated marking schemes for ip traceback," in *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2, pp. 878–886, IEEE, 2001.

[16] M. Caesar and J. Rexford, "Bgp routing policies in isp networks," *Network, IEEE*, vol. 19, no. 6, pp. 5–11, 2005.

[17] J. Van der Merwe, A. Cepleanu, K. D'Souza, B. Freeman, A. Greenberg, D. Knight, R. McMillan, D. Moloney, J. Mulligan, H. Nguyen, *et al.*, "Dynamic connectivity management with an intelligent route service control point," in *Proceedings of the 2006 SIGCOMM workshop on Internet network management*, pp. 29–34, ACM, 2006.

[18] J. Snijders, "Ddos damage control cheap and effective," 2005.

[19] L. Feinstein, D. Schnackenberg, R. Balupari, and D. Kindred, "Statistical approaches to ddos attack detection and response," in *DARPA Information Survivability Conference and Exposition, 2003. Proceedings*, vol. 1, pp. 303–314, IEEE, 2003.

[20] "Slow websites cost retailers billions." `http://mashable.com/2012/11/22/slow-websites/`. Accessed: 2014-10-24.

[21] J. J. Santanna and A. Sperotto, "Characterizing and mitigating the ddos-as-a-service phenomenon," in *Monitoring and Securing Virtualized Networks and Services*, pp. 74–78, Springer, 2014.

[22] M. Vizváry and J. Vykopal, "Future of ddos attacks mitigation in software defined networks," in *Monitoring and Securing Virtualized Networks and Services*, pp. 123–127, Springer, 2014.