UNIVERSITY OF AMSTERDAM

Faculty of Science
MSc. System and Network Engineering

Trusted Networks Initiative to combat DDoS attacks

Jeroen van Kessel
jeroen.vankessel@os3.nl

Alexandros Stavroulakis
alexandros.stavroulakis@os3.nl

April 12, 2015

Abstract

This project consists of a theoretical research, which evaluates the feasibility of the Trusted Networks Initiative from an academic perspective. Initially, we discuss the severity of Distributed Denial of Service Attacks in nowadays world, describe different types of DDoS attacks and the way they work. Secondly, we explain the advantages and disadvantages of different mitigation solutions which help to defend against DDoS attacks. Furthermore, we describe the idea behind the Trusted Networks Initiative, provide insight on the roles of the participants, examine the policies, perform a technical analysis and set different scenarios of attacks and its usage. Finally, we carefully analyze the ethical concerns and risk of such an implementation in real world situations, while also providing some helpful recommendations.

Keywords: Network Security, DDoS Mitigation, Trusted Networks Initiative, TNI

Acknowledgements

We would like to thank our supervisor, Marc Gauw, for his oversight and assistance to our research.

Furthermore, we would like to thank the following parties for providing us with useful information to conduct this research: NLnet, The Hague Security Delta, XS4ALL, ASP4ALL, AMS-IX, NL-ix, SIDN Labs, SURFnet, Bancaire Sector, Ministry of Justice, KPN, UPC and Ziggo.

Contents

Chapter 1

Introduction

1.1 Prologue

In a world where technology is constantly evolving, as the number of Information Technology and Internet applications expands considerably on a global scale, end-users, governments and businesses become more dependent on IT and the Internet than ever before. The interconnection between these parties involves several risks. One of those risks is a digital disruption attack, also known as Denial of Service (DoS) attack. As computational power, bandwidth and network capability keeps increasing, Distributed Denial of Service (DDoS) attacks have also become stronger and more effective[1].

The potential and actual impact of computer and network security threats has become clearer due to a number of highly publicized incidents in the course of the last few years. These threats may not only disrupt our digital infrastructure, but also compromise the integrity, availability and confidentiality of the information we document, analyze and exchange in the digital domain, affecting not only large companies but also end-users[2]. As the cost of these attacks rises, providers, businesses, and governments must respond to protect their investments, revenue, services and customers[3].

A new concept is being developed in the Netherlands, called 'Trusted Networks Initiative' (TNI) aimed to combat these DDoS attacks. This initiative characterizes networks as Trusted and Untrusted, in order to ascertain the validity of internet traffic. Trusted networks are members of this initiative, who agree on specified policies about their infrastructure and network behaviour in order to peer with each other. A member can be any entity with a valid Autonomous System Number that agrees to act according to these policies. All other entities, which are not members of the initiative, are considered Untrusted networks.

The intention of TNI is to help protect against DDoS attacks by allowing 'Trusted Peering' between Trusted parties and temporarily isolating the traffic from Untrusted networks. Trusted Peering is the act of peering between the members of this initiative, also known as Trusted Networks. This paper researches the feasibility of the Trusted Networks Initiative in protecting hosts and networks from large and/or long-lasting DDoS attacks.

The remainder of this chapter concludes the introduction, while Chapter 2 discusses the Research Question of this project and the approach taken to conduct this research. Chapter 3 provides background information on DDoS attacks and describes their causes and severity, while also explaining their different types. Next, Chapter 4 lists the current mitigation solutions, analyzes their manner of operation and examines their advantages and disadvantages. Chapter 5 elaborates on this new initiative, focusing on its concept, policies, technical details and how it would behave in real world situations. Chapters 6 and 7 provide several recommendations for the future development of this initiative and the conclusions of the research, depicting whether or not such a concept is a feasible solution. Lastly, Chapter 8 mentions possible future research on the topic.

1.2 Related Work

The Trusted Networks Initiative, as mentioned above, is a new concept which was initiated in early 2014 and has just passed its test phase. Therefore, no prior research has been conducted and no related work exists on the subject to be used and referenced. From the members of the initiative, only the Policy documentation exists and a minimum form of Proof-of-Concept has been developed[36], on which the technical analysis of the initiative will be based. Similar initiatives or projects developed or already in use, do exist and are described further in the Section 5.5.

The subject of Distributed Denial of Service attacks has been researched multiple times in the past. Furthermore, research papers and data-sheets are constantly published by security firms which provide their customers with solutions against these attacks. Information based on these articles is used in the background research on DDoS attacks and is referenced accordingly when used.

Chapter 2

Research Question

The primary research question of this project is stated below:

Is the 'Trusted Networks Initiative' a feasible additional solution in protecting hosts and networks from large and/or long-lasting DDoS attacks?

To fully explore the Trusted Networks Initiative and form a conclusion about its feasibility, the following underlying questions have been formulated:

General/Policy Aspect

- Is the 'Trusted Networks Initiative' feasible in the Dutch national landscape considering the involved networks, hosts and internet exchanges?
- Can the 'Trusted Networks Initiative' expand to international networks?
- How does the 'Trusted Networks Initiative' relate to initiatives like 'Routing Manifesto' (ISOC) and 'Fenix ' (Czech Republic) and can the projects align?
- Are the policies of the 'Trusted Networks Initiative' well defined?

Technical Aspect

- Is a concept of Trusted Routing between hosts and networks anyhow a feasible solution that should help mitigate DDoS attacks?
- What are the other technical solutions to mitigate DDoS attacks?
- Is there a serious risk that DDoS attacks, indeed, become too large and/or too prolonged to mitigate?
- Can the 'Trusted Networks Initiative' easily be combined with the other DDoS mitigation solutions?

End-User Aspect

- How will end-users experience the availability of 'critical websites' (e.g. banks, government, energy) when a DDoS attack takes place while the 'Trusted Networks Initiative' is active?
- How would this compare to a situation where none of their customers can reach a critical website anymore?

2.1 Approach

During this research, no facilities were provided or made available in order to perform tests on the Proof-of-Concept of the Trusted Networks Initiative, which could simulate attack situations in an appropriate environment. Therefore, in order to answer the antecedently stated questions the initial steps were to come in contact with the members of this initiative and discuss their motivation for joining and supporting this initiative. Along with what benefits their participation has both for them and the initiative as a whole. As the Trusted Networks Initiative is mostly a collaboration based on policies, which attempt to ensure that the members form a web of trust, special attention was paid on their effect in the initiative's operation. By analyzing these policies and the Proof-of-Concept, and collecting technical information about its inner workings, different technical scenarios were designed and theorized to detect its effectiveness.

The intention of this research is to theoretically evaluate this new concept and produce a detailed conclusion on its feasibility. This will be achieved by attempting to answer the aforementioned research questions and by giving, in the end, a shape to this initiative, while also trying to provide some helpful recommendations regarding its policies and method of operation. The key employees of the members of TNI were interviewed during the research. Since TNI is still a concept in their minds and has not, yet reached its implementation phase, the thoughts of these members were heard and collected in order to be put together in this report. Many of them do not share the same idea about the final form of the initiative. Therefore, this paper tries to align the actual implementation form of the initiative and investigate which way would be more beneficial.

Due to the limited resources, which were made available throughout the course of this research – beta version of the Policy documentation, short PoC – the chosen approach was to interview the initiators and shape their visions of this concept into a single cohesive document. This way a better in-depth understanding of this initiative can be provided.

Chapter 3

Distributed Denial of Service Attacks

3.1 Background Information

Distributed Denial of Service Attacks (DDoS) are a form of a computer network attack that first originated in the beginning of the millennium and have evolved into a strong weapon in the attackers' arsenal. One of the first recorded DDoS attack incidents was in 1999, with the University of Minnesota's IRC server as the target, an attack, which lasted for several days[4]. In hindsight, this could be considered an omen of what were to follow; in February of 2000, a teenager from Quebec would launch a DDoS attack, which managed to shut down the – at the time – top used search engine, Yahoo!, for almost one hour and also targeted the websites of eBay, CNN, Amazon and Dell among others, something which had great compound losses as a result[5]. In the last few years, these attacks have grown both in severity and frequency. Their duration varies from a few minutes to multiple days or longer and the reasons for these types of attacks vary from pure curiosity or fun to illegal profit gaining or ideological 'hacktivism'.

What these attacks need in order to succeed, is volume. First and foremost, they work in a distributed model. The attacker behind a DDoS attack needs a significantly large number of devices with internet access which can be remotely controlled and when a command is issued, the devices can execute it and begin the attack. Those are called 'zombie' computers which become part of a 'botnet', a network of slave computers which can be used by the attacker. When the botnet begins the attack – which varies in type – each one of those devices initiates a packet flow towards a chosen target which is too large for the victim to handle, causing, among others, either network congestion or CPU overloads and renders the target service unavailable. Therefore, other end-users will be unable to access the offered services.

Of course there are many variants besides the large quantity of connections that need to be taken into consideration when defending against DDoS attacks. For a successful DDoS attack, the point of origin, in the majority of times, should not be detectable. Therefore, the source IP addresses can be 'spoofed', meaning they appear to be not the originals. However, a DDoS attack with a legitimate IP range can also cause significant problems. Another attack can be a 'reflection' type of attack, where a legitimate request is sent to a publicly available server, with the source IP address being the one of the victim, which has as a result the public server to send the response to the victim. This means that pinpointing the attacker's location is rather impossible. By having the ability to fake or conceal their point of origin, an attacker can remain hidden and if the source of traffic cannot be filtered and identified, the attack cannot be stopped at its source.

The different types of DDoS attacks and the current ways to mitigate them shall be discussed further.

3.2 Motivations and Causes of DDoS Attacks

The purposes of DDoS attacks vary. A small percentage of them can be accidental, caused by a badly configured system or as a demonstration to potential customers of DDoS protection solutions. However, there is often a personal intention behind the majority of these attacks. Some can be used as a diversion for attackers who want to try and steal information from certain systems. Or for financial market manipulation and/or taking out competitors. Or even fame gaining, because hackers might want to boast that they managed to successfully attack a well known target. Moreover, online gaming, gambling and social network related reasons are also motives behind such attacks. One of the biggest motivations behind DDoS attacks is ideological 'hacktivism'. Hacktivism is mainly driven by political and ideological disputes[7].

The majority of these attacks have their origin outside of Europe and especially outside of the Netherlands[8]. Figure 3.1[1] shows the top 10 countries of origin of DDoS attacks.

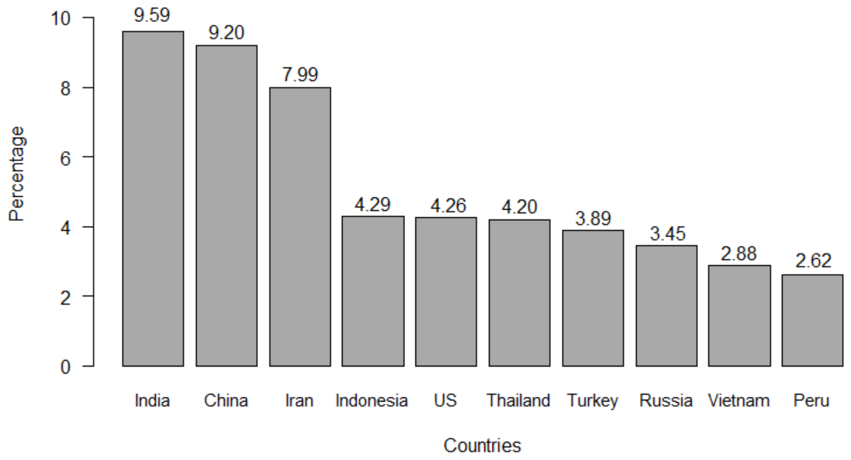

Figure 3.1: Top 10 countries of origin of DDoS attacks (in %)

Those are the countries from which the resources of the attackers originated. These resources, which are part of a botnet, can be defined as personal computers, routers, servers, whole data centers, or open resolvers which can be accessible from the internet. Botnets can be everywhere and this is due to the fact that regular end-users with infected computers are everywhere and well built botnets are difficult to detect[43]. This means that most resources can be found where computers are common and little to no security measures are taken.

An important issue is awareness and that not many Access Providers have actually taken into account the fact that they are a part of the problem[9]. According to the 'MIT ANA Spoofer

[1] Incapsula Top 10 DDoS Attack Trends of 2014

Project', close to eighty percent of the Providers has reached ingress filtering deployment[10]. However, the remaining twenty percent is rather large, which results in a great influence in withholding the anonymity of the attackers.

Illegitimate IP addresses help to keep the location of the attackers secret. Lists of IP ranges or DNS servers which can be used in these attacks are well known among hacker groups and are being exchanged in online black markets or on the Darknet[11]. This, of course, leads to offering DDoS attacks as a form of a service (DDoSaaS) and for the aforementioned motives, by paying the desired price, one can request a DDoS attack on a certain target.

This, does not mean that the only targets of these attacks are large companies such as banks, media companies or governmental services. DDoS attacks have a broad spectrum of targets whether these are regular end-users at home, or infrastructures like Web servers, File Servers, Trading Platforms, VoIP Systems, E-mail Servers, DNS Servers, Gaming servers and Digital Banking platforms. Anyone can become a target of a DDoS attack.

3.3 Severity

As already mentioned, with the passing of time and the evolution of technology, DDoS attacks have grown stronger, meaning the mitigating solutions are becoming even more expensive, and certainly indispensable. The Dutch National Cyber Security Center reported an exponential increase in volume and size of DDoS attacks over the last couple of years. This increase in bandwidth resulted in unreachable websites and collateral damage within the IT-infrastructure; especially the financial sector and governmental departments, which were the victims of these attacks[6].

Furthermore, large-scale DDoS attacks are getting more frequent. Especially DDoS attacks of more than 10 Gbps are getting more common[12]. On-premise mitigation devices cannot withstand attacks of this volume if the upstream bandwidth capability is exceeded. In 2013, security firm Arbor measured a DDoS attack of 309 Gbps[7] (figure 3.2), while Neustar reported an amplification attack of 400 Gbps[13]. DDoS attacks of this scale require a new solution in order to mitigate the DDoS attacks of the future, which will most probably increase in bandwidth.

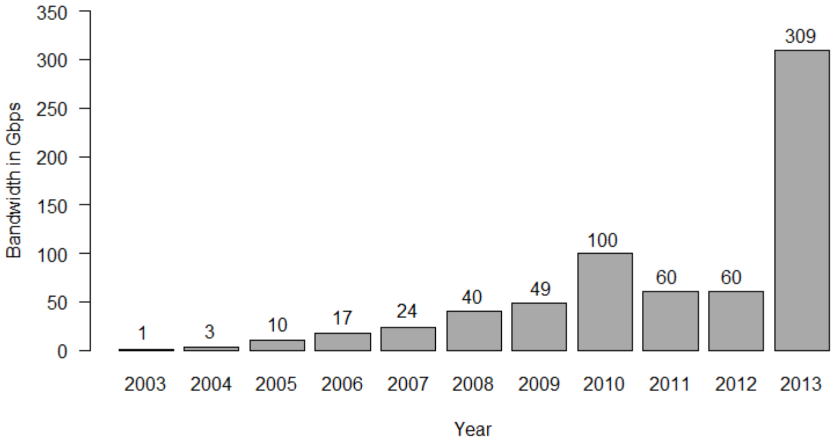

Figure 3.2: Size of largest reported DDoS attack (in Gbps)

DDoS attacks are becoming a more and more popular weapon of choice since they provide anonymity to the malicious initiator. The attacks have known a significant rise in frequency the last few years and according to Verisign, more than 10,000 attacks occur each day[12]. The intervals between these attacks are random, but a large number of them repeat themselves targeting the same victims over and over, which shows the attackers really do take advantage of every small vulnerability they discover[13].

While the bandwidth and frequency increase, the duration of these attacks seems to be decreasing, with the actual intention being to be short and effective. These short attacks appear with a rate of eighty-eight percent for a duration of up to sixty minutes[7]. The final goal of the attackers is to cause failure of operation to a certain system, however their attacks can last longer if their ulterior motive is to prove a point or simply show off their abilities[12].

3.3.1 Spamhaus Case

The infamous Spamhaus case of 2013 is a realistic DDoS scenario for future attacks in nowadays society. Spamhaus is a non-profit organization which performs spam-filtering by maintaining a block-list of IP addresses and recommending to discard e-mails originating from this block-list. This block-list is globally used by Hosting Providers, enterprises, governments and many other organizations[14]. E-mail spammers have always been in disagreements with Spamhaus filtering their illegitimate e-mails. This disagreement resulted in a DDoS attack against Spamhaus. A Dutch Access Provider (ISP) hosted one of the Spamhaus block-list servers in their data center, which resulted in a collateral DDoS attack. Figure 3.3[2] visualizes the bandwidth sample across a number of the routers in front of these servers[15].

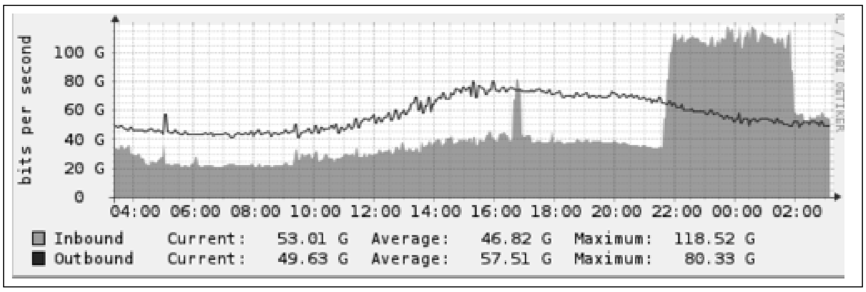

Figure 3.3: In-bound and out-bound traffic of the routers during the DDoS attack on Spamhaus, March 20 2013

The line represents the out-bound traffic while the static area represents the in-bound requests. The attack started with a bandwidth of 10 Gbps and reached 120 Gbps by the time the Access Provider realized the attack was too large to handle. Their final decision was to take the Spamhaus servers offline.

As can be seen in figure 3.2, it is evident that DDoS attacks have seen exponential growth in bandwidth and frequency, while also managing to exceed the current expectations.

[2] M. Prince, CloudFlare, The DDoS That Knocked Spamhaus Offline, 2013

3.4 Distributed Denial of Service Attack Types

In essence there are two basic types of DDoS attacks - volumetric and application layer attack [40] [41]. Volumetric attacks are flooding attacks whose aim is to saturate and consume the network bandwidth and infrastructure of the target. Application Layer attacks tend to use relatively less bandwidth and to be harder to detect; their targets are applications and/or services where they gradually consume and exhaust resources [7].

3.4.1 Volumetric attacks

Volumetric attacks, otherwise known as Layer 3/4 attacks, from the OSI model's Transport and Network layers, can be split into two categories, Flood Attacks and Amplification Attacks [42] [40].

Flood DDoS attacks

The idea behind Flood DDoS attacks is to send extremely large amounts of traffic to the target with the goal being to overburden its network [40]. This way, legitimate traffic cannot reach its destination, or responds so slowly and is then rendered essentially unavailable. Such attacks usually lead to a server overload or crashes[16]. The most common Flood attacks use, among others, TCP SYN/ACK, UDP, ICMP packets and recently popularity has grown in the use of SSDP packets[12] [42].

Amplification Attacks

Amplification DDoS attacks try to communicate with a specific broadcast IP address which results to traffic being sent by the entire subnet, reachable from this address, to the target of the attack[39]. Throughout the broadcast address, the traffic is amplified and reflected towards the victim which then of course affects their bandwidth [42]. A shortened example of this method is shown below where a simple DNS dig query is used to obtain a relatively large response from the Dutch governmental authentication website 'DigiD'; a query with the size of just 64 Bytes, results to an answer of approximately 3 Kilobytes. Another example that is part of a DNS Amplification DDoS attack is a 'Smurf DDoS attack', in which the attacker uses the victim's spoofed IP address, resulting to the amplified traffic being sent back to the victim.

DNS Lookup

```
\$ dig ANY digid.nl +edns=0

; <<>> DiG 9.9.5-3-Ubuntu <<>> ANY digid.nl +edns=0
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8822
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 27, AUTHORITY: 5, ADDITIONAL: 8

;; QUESTION SECTION:
;digid.nl.                          IN        ANY

digid.nl.               600         IN        RRSIG        A 8 2 600 20150415
07293620141216062936 35550 digid.nl.
hIkrfQLxGZChB1ezeweXbY6byZh4r7xvCRU/y1m2XXKTKydKiJf+0ap3 mGY0unc4mVxBRRItLoiMw
gH/Qc2onYO/7mY1CeSc1Tz5EK8S8jRduc/3 kVlaAsn9ZB//Xw2O8IXUDASv0qbyvTVlL9dvlJf7OI
o9pilrWkJXKEtW wdk=

[...]

;; Query time: 21 msec
;; SERVER: xxx.xxx.xx.xx#53(xxx.xxx.xx.xx)
;; MSG SIZE  rcvd: 3128
```

3.4.2 Application Layer Attacks

Application Layer attacks, also known as Layer 7 attacks, from the OSI model, tend to imitate human behaviour and target specific parts of systems which provide user interaction, which makes them harder to detect [40]. In essence, rather than trying to saturate an entire network connection or congest a server with enormous amounts of traffic, these attacks focus on a rather small number of applications or operating systems such as Windows, Content Management Systems, Webmail applications, or taking advantage of the way specific systems work, e.g. by exploiting Apache's thread creation for each new connection and sending packets in slow speeds through multiple connections[42]. This type of DDoS attacks is currently on the rise[17]. Examples of those attacks are Slowloris or Slow Read attacks, VoIP Amplification attacks and HTTP flooding attacks [18].

3.5 Conclusions

This chapter focused on the motivation and causes of DDoS attacks and their different types. As with all security issues, preventing is preferred over detecting and mitigation. Organizations must protect weak points of their infrastructure. The different attack types were described and categorized according to their targeting layers. The fact that there are more than one target points for DDoS attacks – Layer 3/4, 7 – reveals that only one approach to the prevention of DDoS attacks is not enough. A multilayered prevention and mitigation approach is required to prevent both Volumetric and Application Layer attacks. Providers should implement measures against spoofed IP addresses. These measures should result in limiting the origin of malicious packets. Furthermore, a forensics investigation should be more feasible. Volumetric attacks need to be mitigated at the Access or Content Providers side before they overwhelm the local network or the security infrastructure. Application layer attacks need to be defended against in the target's data center or on-premises[40].

If compliance in terms of surveillance or by law is not urgent, protecting personal data should always be preferred. Less availability might be less interesting compared to integrity. Implementing a new solution, which makes organizations more available, might result in less integrity. Higher security in terms of integrity can create weaker links in terms of availability. Organizations should find a striking balance to protect availability, integrity and confidentiality.

Chapter 4

Distributed Denial of Service Mitigation Solutions

Due to the amount of different DDoS targets, various types of attacks and their severity, it is necessary to have methods to mitigate these attacks and help secure the availability of critical services. Nowadays, there are different types of popular techniques, which attempt to accomplish this. However, this is not always performed successfully[40]. Non-profit and commercial mitigation solutions help organizations to defend against these types of attacks. These solutions can provide a high success rate, but they can also prove to be very costly, depending on the size of the attack.

4.1 Popular DDoS Defensive Techniques

The current most popular defensive techniques involve 'Traffic Blackholing', which is a way of discarding all incoming traffic towards the target in order to save the Hosting Provider's network from saturation[3] [45]. However, by discarding traffic in its entirety, legitimate traffic can also be dismissed. Access Control Lists of routers is another choice of protection. This is successful in filtering already known attacks by examining the protocols used. However, DDoS attacks are becoming more sophisticated and use valid protocols, rendering this filtering unsuccessful when it comes to SYN or SYN/ACK DDoS attacks – a form of flooding attack, which involves sending Sync and Acknowledgment packages at a high rate[42]. Another way of using routers as a form of protection is to use Unicast Reverse Path Forwarding (uRPF), which can be used to block IP addresses outside of the target's subnet[18]. Nonetheless, if an attacker uses spoofed IP addresses from the same subnet, there is little to nothing that can be done. Also, legitimate end-user traffic is blocked and thus the DDoS attack succeeds.

According to Cisco, another popular opinion is that firewalls and Intrusion Detection Systems (IDS) are an adequate form of protection against DDoS attacks[3]. Firewalls are used in-line and because of that, attackers target them for their low session handling abilities. Usually, they do not filter spoofed traffic and they can also be used to reject traffic from certain protocols. However, the attacking side can still use valid protocols during a DDoS attack. The same applies to IDSs, they can provide excellent detection for Application Layer attacks but not against valid protocols. Moreover, as the term suggests, IDSs only function as a detection mechanism[18].

4.2 Commercial Solutions

Contrarily, the commercial solutions have a high success rate in mitigation. The companies offer different types of plans according to the clients' needs and budget. Services such as prevention, monitoring and traffic handling can be provided. There are mainly two different types of traffic handling which are correlated to the layers of attack, Layer 3/4 and Layer 7 solutions[8].

Layer 3/4 mitigation techniques are based on Border Gateway Protocol (BGP) IP address range swings. This technique is also called 'Off-Ramping'[25]. The target of the DDoS attack can decide to stop announcing their IP address range to the global Internet and in turn, that particular company announces it for them, meaning that they receive all of the traffic intended for the client, whether it is malicious or not. This operates in a distributed model, with data centers across the world, traffic is sent to the nearest, in order to limit the amount of latency[41]. The traffic is then 'washed', as it goes through special purpose built appliances to filter illegitimate traffic out with the use of specific algorithms. Once the traffic is 'washed' it is rerouted back to the client (On-Ramping) over a Generic Routing Encapsulation (GRE) tunnel[25].

Layer 7 mitigation techniques also function in a distributed model. By having multiple data centers at different Internet Exchanges, the clients can point the DNS entry of their websites to these companies who in return, handle all the requests where each packet is then inspected[44]. Thereupon, based on the signatures, illegitimate traffic can be detected and discarded. Next, legitimate traffic is sent back to end-users' browsers based on their geographical location[26].

Figure 4.1 shows essentially how these solutions work. Due to their immense amount of available bandwidth, both legitimate and malicious traffic is accepted. The traffic is then washed using algorithms to examine which packet protocols are used. At the end, they discard the DDoS traffic and send the legitimate traffic to the Critical Services Infrastructure and back to the end-users.

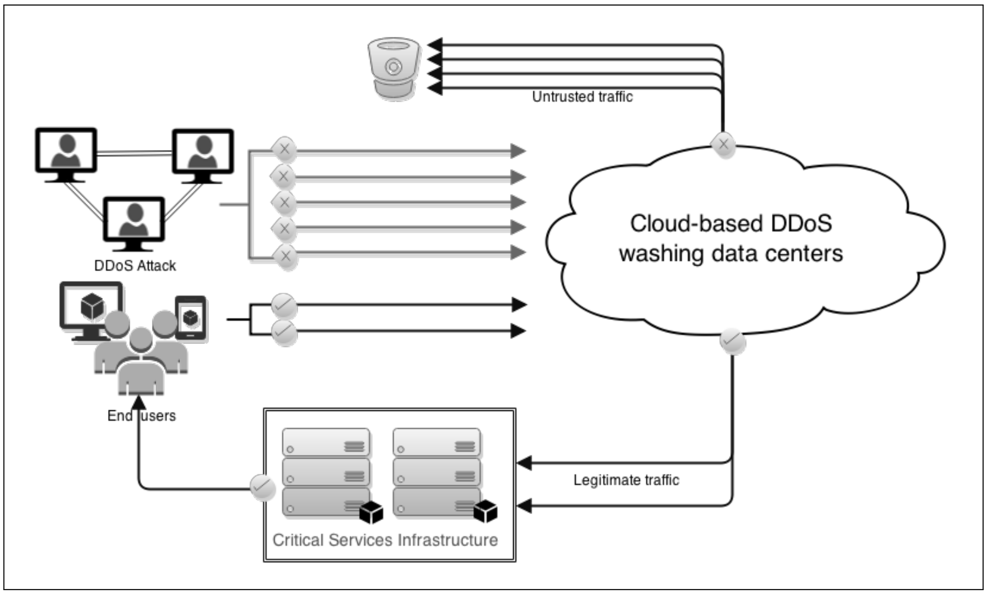
Untrusted traffic
DDoS Attack
Cloud-based DDoS washing data centers
End users
Legitimate traffic
Critical Services Infrastructure

Figure 4.1: Overview of Cloud-based DDoS Mitigation techniques

4.3 National Anti-DDoS Wash Street

The Dutch Association of Internet Providers 'Nationale Beheersorganisatie Internet Providers' (NBIP) initiated the 'National Anti-DDoS Wash street' (NaWas) [23]. NaWas is a non-profit organization used in collaboration by the Dutch Access and Content Providers. It is a shared facility of approximately 960 Gbps bandwidth in total and can mitigate a large-scale DDoS attack. Working in a similar way as explained above, the participants can swing their BGP announcement(s) to the NaWas data center in order to 'wash' the DDoS attack traffic. It is intended to be used as a deep-packet washing data center by the Dutch Access and Content Providers to mitigate DDoS attacks as a cheaper solution than using a commercial one[23].

4.4 AbuseHUB

Another Dutch initiative is AbuseHUB, where Dutch telecommunication companies and Internet Providers collaborate in order to collect and analyze reports of botnets[43]. These reports are sent back to the Internet Providers, who then have an overview of infections in their network and can take the appropriate actions in removing them, after informing their customers. However, this initiative is not meant to be a solution, but an after-measure which can help reduce related attacks in the future with the intention to sanitize the Providers' networks and educate the end-users about their online activity and computer usage[24].

4.5 Disadvantages of Commercial Solutions

Commercial mitigation techniques prove to be quite efficient, however, they are not perfect and at times – also depending on the customer – can be quite expensive as well[25]. First and foremost, the algorithms used to 'wash' traffic are not flawless, which means that along DDoS traffic, sometimes legitimate traffic is discarded. These mitigation solutions offer DDoS detection and contact the client when they detect a significant rise in their traffic and ask if they should take measures[44]. Before the detection has provided results and the BGP swing of the IP range has taken place, up to thirty minutes can pass, during which the victim is under attack and unable to react[19].

Application Layer solutions have a different disadvantage. Because of their distributed model, which can have a replica of the client's web service in any of their data centers, it is rather unsafe to use this for services that implement SSL, due to the fact that the Private Keys would need to be shared. For small customers, this might not be an issue, however, when the customers are banks and/or critical governmental services, the privacy issues which arise are far greater. Recently a solution to this problem was developed[20] which tries to provide this mitigation solution without the sharing of Private Keys. Nevertheless, the matter of privacy still remains when bringing a third party into the equation. End-users of critical services will be forced to trust their financial and private data to be handled by an entity which provides them with no transparency.

Lastly, these companies offer contracts based on bandwidth. If a customer, for example, chooses a contract of mitigating attacks of up to 40 Gbps and is attacked with a larger DDoS attack than the one stated in the contract, the prices of mitigation increase excessively. According to security firm Imperva, sixty percent of US companies experienced DDoS attacks during 2013. DDoS mitigation solutions can cost between $5,000 and up to and over $100,000

US dollars per hour[21]. Furthermore, these figures do not include possible reputation damage and/or customer satisfaction. The consequences of a DDoS attack can thus be disastrous to any organization or company.

4.6 Conclusions

As explained above, commercial mitigation solutions are sophisticated techniques, but depending on the attack and the target, they may not always be sufficient. Ultimately, their goal is to 'wash' the traffic and procure evidence which can identify the attacker and can lead to a criminal court case. However, this is not always possible, since as recent history has shown, and was mentioned in the previous chapter, the attackers always manage to find new techniques of attacking. For that reason, there is a need in variety of defenses. The attackers' bandwidth capacity is inherent to the continuous growth of DDoS attacks. Therefore, adding more bandwidth to DDoS mitigation solutions is not the answer in the long term. Both the customer side and the DDoS solution providing side should be ready for the DDoS attack of the future, which will be too large to mitigate and will last long enough to threaten with bankruptcy even the most profitable companies. The current solutions exist to defend against the current DDoS attacks and not the ones which might occur in the near future.

In addition, the amount of time until the mitigation starts providing results may not be quick enough. Moreover, being inaccessible for thirty minutes can prove to be quite costly for organizations of any economic scale[19]. This reveals the need for an alternative solution for both small and large organizations. This is the main reason behind the concept of the Trusted Networks Initiative. It comes as a security building block to help provide better results along with the commercial solutions. Whether it succeeds or not in its aim, will be evaluated in the following chapter.

Chapter 5

Trusted Networks Initiative

5.1 Concept

The Trusted Networks Initiative is a non-profit, internationally oriented initiative, designed as a last resort communication channel for its members in the event of a large scale DDoS attack[31]. By defining trusted parties, known as 'Trusted Networks', it intends to form a high level of confidence and security and ensure the communication of the participants and the availability of their provided services. As briefly mentioned in Chapters 1 and 2, the main idea of TNI is trust, therefore its operation is mostly based on policies and the way in which its members act. It separates its members, the Trusted Networks that agree to the policies, from the Untrusted, which are entities not participating in the initiative.

Trusted Networks can re-actively choose to use the Trusted Routing service as a last resort measure in case of a DDoS attack. The goal of this initiative is to ensure that its participants can maintain and uphold their critical connectivity between important applications and the 'local' access networks during such attacks, independently from traffic with the untrusted part of the Internet[32]. Figure 5.1 shows a high level overview of the Trusted Networks Initiative.

In the figure, the end-users – and customers – of an Access Provider are depicted in the bottom left corner, and want to connect to a Critical Service behind a Content Provider at the bottom right corner. That Critical Service is also the target of a DDoS attack issued by an attacker shown at the top of the figure. The end-users can access the Critical Service in two different ways; over the Internet, where their traffic goes from the Access Provider, through the Internet Exchange Point (IXP), to the content Provider and then reaches the Service. Or, from the Access Provider, through the TNI overlay network, to the Content Provider and finally to the targeted Service. While the attack traffic can only reach its target through over the public Internet, and not through the TNI network.

'TN' stands for Trusted Network, which, as stated above, is a qualified network of this initiative that complies to its rules. 'Trusted Routing VLAN' is the overlay network of the Trusted Networks Initiative, over which its members can communicate. The use of the Route Server, which will be discussed further in Sections 5.3 and 5.4, is to maintain the peering connections between the participants.

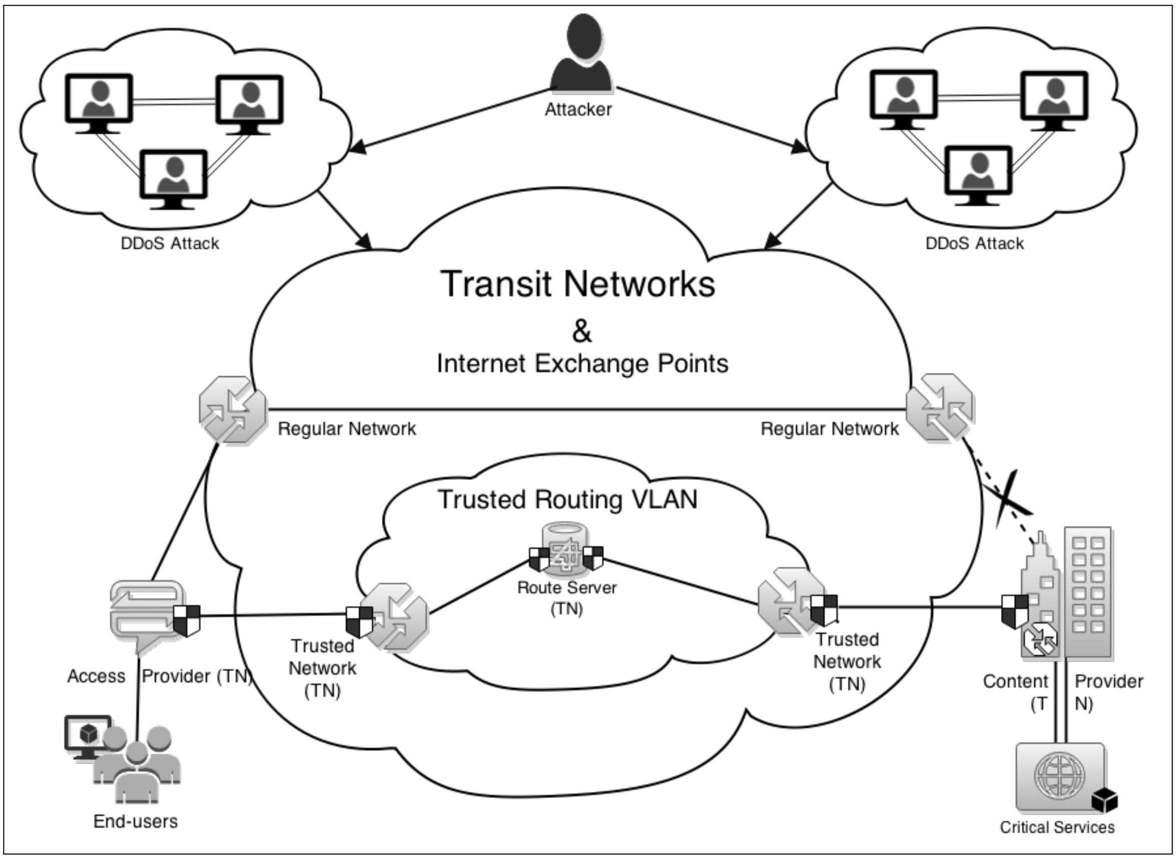

Figure 5.1: High-level overview of the Trusted Networks Initiative Concept

By temporarily disabling the Transit connection, ergo the access to the public Internet, end-users of a Trusted Network Access Provider will be able to maintain connectivity to critical services of possible DDoS attack targets, while the targets can preserve their service's availability and functionality even under long lasting attacks. The members of the initiative believe that by segregating approximately 98,5 percent of the DDoS attack traffic, the remaining 1,5 percent is from within the trusted networks. This will result in a significant decrease in DDoS traffic, which is manageable by on-premise mitigation solutions. End-users will, therefore, be able to access their critical services in case of a major DDoS attack. Moreover, with the use of ingress filtering techniques, detecting the source of the attack becomes easier, with the benefit of this being the discovery of the originators of these attacks.

The policies are oriented towards DDoS protection and safety and a way to minimize abuse within the Trusted Networks.

5.2 Policies

The Trusted Networks Initiative is created with the aim of providing, at minimum, a last resort interconnection to other Trusted Networks, when one of the collaborating entities becomes a target of a DDoS attack[31]. By creating a safety ensured group of Access and Content Providers and employing best practices as policies, participants should be able to monitor malicious traffic.

The Hague Security Delta (HSD), the largest security cluster in Europe, gave this initiative a platform by chairing this project[33]. Furthermore, geographically, the Netherlands is the digital gateway to Europe. The Trusted Networks Initiative's policies describe the qualification rules to which a participant should comply, in order to become a member.

Any organization can become a Trusted Network by providing its Autonomous System Number(s) (ASN), which must be connected to either the Amsterdam Internet Exchange (AMS-IX) or the Neutral Internet Exchange (NL-ix). Furthermore, the organization's Network Operations Center (NOC) must be functioning in an incessant basis with at least two IT-professionals. These actions should result in internal incident resolution procedures within the organization and in a capable Computer Emergency Response Team (CERT) or Computer Security Incident Response Team (CSIRT). Moreover, cooperation with the authorities is an important aspect of this initiative in order to prevent the same DDoS attack taking place in the future. This cooperation can result in a forensic investigation to trace the cause and possible suspect(s) of the DDoS attack.

From a technical aspect, IP source address spoofing measurements must be implemented according to RFC 2827[27] combined with BGP4-router(s) to route traffic with other participants [34]. RFC 2827/BCP 38 describes ingress traffic filtering to prohibit DDoS attacks which use forged IP addresses to be propagated from 'behind' an Access Provider aggregation point[27]. The urge to implement these best practices resulted in a national implementation within almost all Dutch Access Providers. Additionally, tracing an attacker or infected client should be easier since a legitimate source address must be used.

Furthermore, non-managed, open DNS resolvers must be taken care of, as well as implementing rate limiting against DNS amplification attacks and/or similar attacks. Moreover, the organization must implement a monitoring service on their network in order to actively detect and signal an irregularity in the monitored values. Also, the organization's router's control plane must be protected against undesired or malicious traffic in compliance with RFC 6192 (Protecting the Router Control Plane)[35].

Although traffic monitoring, anti-spoofing and router control plane protection measurements are mandatory, a DDoS attack can still originate from within the Trusted Networks. This means that either a user connected through one of the participating Access Providers can initiate an attack with a member of the initiative being the target; or that an outsider is able to control to a group of infected computers, which reside within the initiative and are part of the network of one of the Trusted entities. If so, a best effort of one hour is reasonable to solve this issue. A participant can also be suspended or expelled when violating the Trusted Networks Policies.

Communication between members will be achieved via mailing list. Each of the participants has the obligation to provide other members with information regarding any significant security incidents, which the initiative is supposed to prevent. The Trusted Networks Initiative is meant to be a collaboration of trust between its participants, which should result in a safer Internet. This is the reason the members of this initiative focus more on its policies and the relationships between the participants.

5.3 Technical Analysis

On the 27th of June 2014, the Trusted Networks Initiative passed its testing phase according to the Proof-of-Concept, on which this analysis and the following scenarios are based[36]. Figure 5.1 is based on this document where members (Access and Content Providers, Internet Exchange) of this initiative came together to examine its functionality by performing tests under regular circumstances and under DDoS attacks.

For this test, two dedicated lines of 1 Gbps were used to connect a test website hosted at a Content Provider with the public Internet – via an Internet Exchange – and with an Access Provider – via Trusted Routing. Also, an available IP range and an unused webserver with two separate IP addresses available to the public internet were used. By attempting to access the website with `ping`, `traceroute` and `wget` tools, the involved parties were able to confirm the fact that both the website's IP addresses are reachable from the Access Provider client through both regular peering and Trusted Routing under regular circumstances. This setup is shown in figure 5.2.

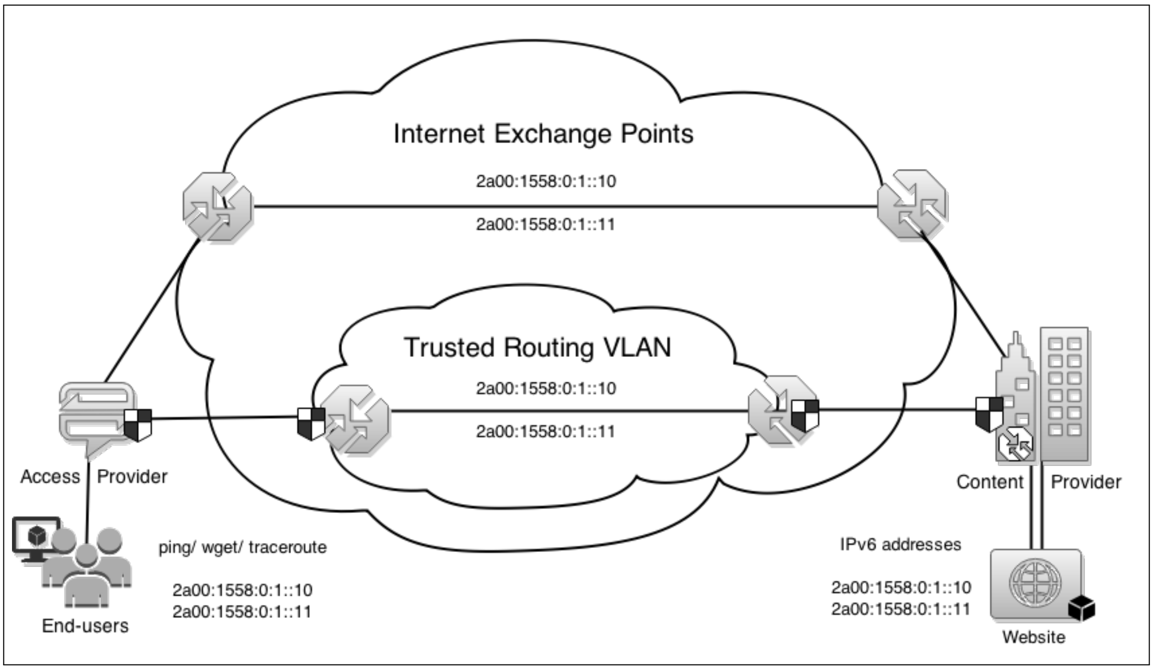

Figure 5.2: TNI Test Setup

For the second part of this test, a DDoS attack was issued from a third party, via the public internet, targeting one of the IP addresses of the website hosted at the Content Provider's side. The same commands as above were used to determine connectivity. During the attack the client at the Access Provider's side was unable to reach either of the IP addresses via the Transit. Then, by performing Remote BGP Blackholing[45] – which means that the attack traffic was discarded before it reached its target – at the Transit connection, the attacked IP address was reachable via Trusted Routing but not via regular peering, while the second IP address was reachable from both. This is shown in figure 5.3.

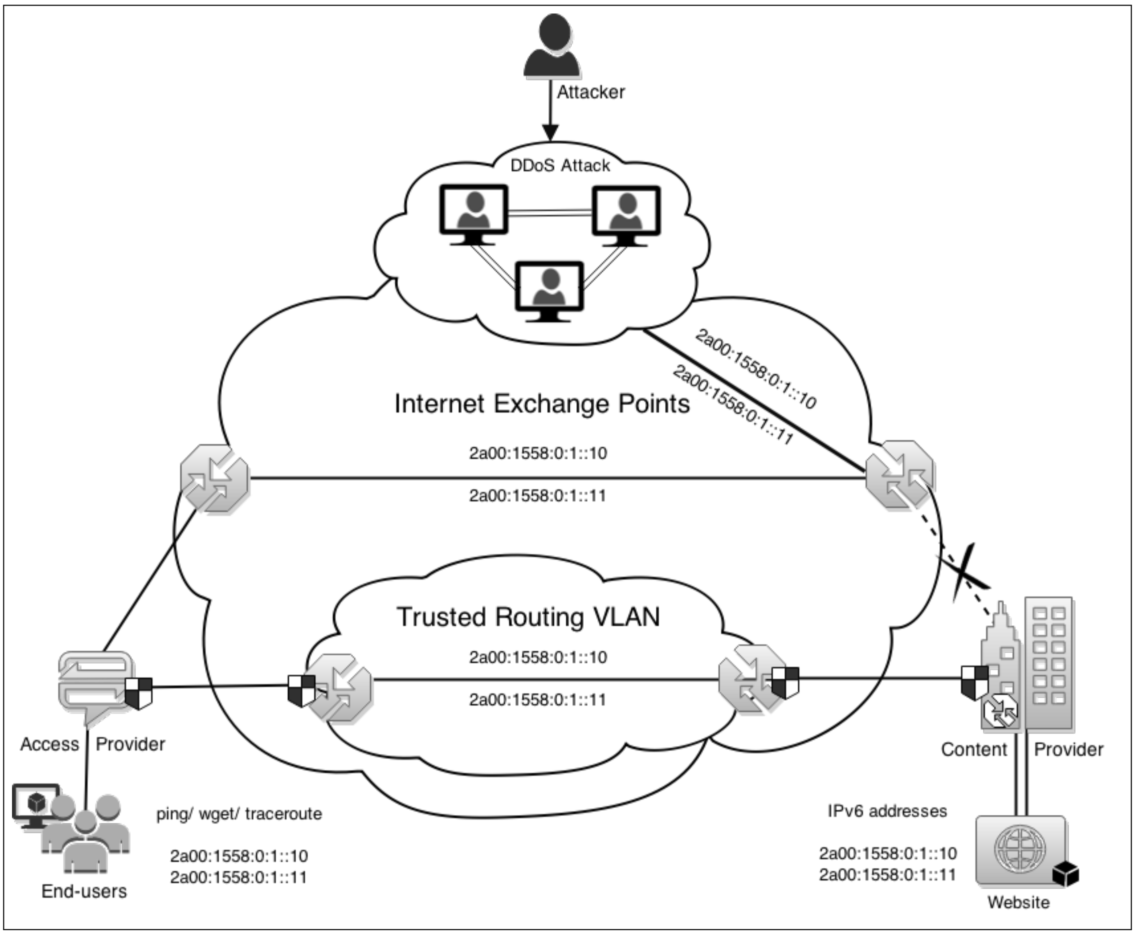

Figure 5.3: DDoS Attack

For the last part of the test, like above, one of the IP addresses of the website was once again under DDoS attack via the public Internet. However, instead of Remote BGP Blackholing[45], the Transit connection was disconnected. As a result, both IP addresses of the website were reachable only via Trusted Routing.

When a heavy DDoS attack occurs, a participant has the option to disable the Transit connection where the attack is coming from. However, doing so should only be temporary. Therefore, detecting whether the attack is of crucial matter. There are two ways to disable the Transit connection. The first option would be to null-route[45] at the recipient's side. That way, the DDoS attack will still be visible by its target. The attack can be monitored by bandwidth and the load of the router's CPU. Further procedures can be taken when the DDoS attack is becoming less intense or disappears.

The second option is to request the upstream Provider to null-route the connection[45]. This can be achieved by configuring the settings within the router to block the IP address. However, monitoring the DDoS attack at the recipient's side is not possible. The upstream Provider could be contacted for statistical analysis based on the connection in order to detect whether the DDoS attack is still active or has ceased its action. By discarding the attack traffic at the upstream provider's side, the devices at the recipient's side – such as firewalls, routers and the service nodes – will not have to cope with the immense compute load, because traffic will not arrive there in the first place.

The initiative aims beyond the scope of Dutch participants. Since no discrimination takes place – according to the claims of its members – any party from separate and distant parts of the world can become a member of the TNI. However, a slight increase in latency can occur when deploying this initiative on a worldwide scale. For now, such an increase should not occur since this initiative and its participants are all located in the Netherlands. By placing a route server at the Internet Exchange Point, participants can easily maintain multiple connections with everyone by building one session. This measure is taken to ensure the scalability of the initiative in the event of it expanding outside of the Netherlands.

However, it is not mandatory to use a route server. Participants can also configure their BGP-4 router(s) to connect to all other participants by AS-number in order to peer with each other. All participants can then independently decide for themselves to accept or refuse traffic from other members of the initiative, in case of a heavy DDoS attack from within.

The technical core of the Trusted Networks Initiative is the Trusted Routing VLAN on BGP level. This, so called, logical BGP peering group will use the already existent infrastructure by implementing this initiative as a logical overlay and route traffic through the Trusted Routing VLAN to the other participants. The Trusted Routing VLAN consists of Access Providers, Content Providers and Internet Exchanges, governmental departments, commercial companies, financial companies, energy selling companies etc. These parties will have to agree on the above policies, on a qualified infrastructure and proper network architecture to route traffic between their closed group of members and to take the appropriate actions when necessary in the least amount of time. This leads us back to the figure 5.4 which shows the high-level overview of this initiative.

5.4 Scenarios

As the intention of TNI is to operate along with the current DDoS mitigation solutions, there is not only one manner in which it aspires to succeed. Two implementation scenarios are possible. The Trusted Routing VLAN can be activated either when a DDoS attack occurs targeting a participant or at all times with the highest priority.

The first scenario, is the intended use of the Trusted Networks Initiative, as a temporary and last resort solution. As previously mentioned in the above subsection, when a member of the initiative is under attack, they can decide to refuse all public traffic and accept only the incoming traffic from the Trusted Routing VLAN. Thus, all other members can communicate

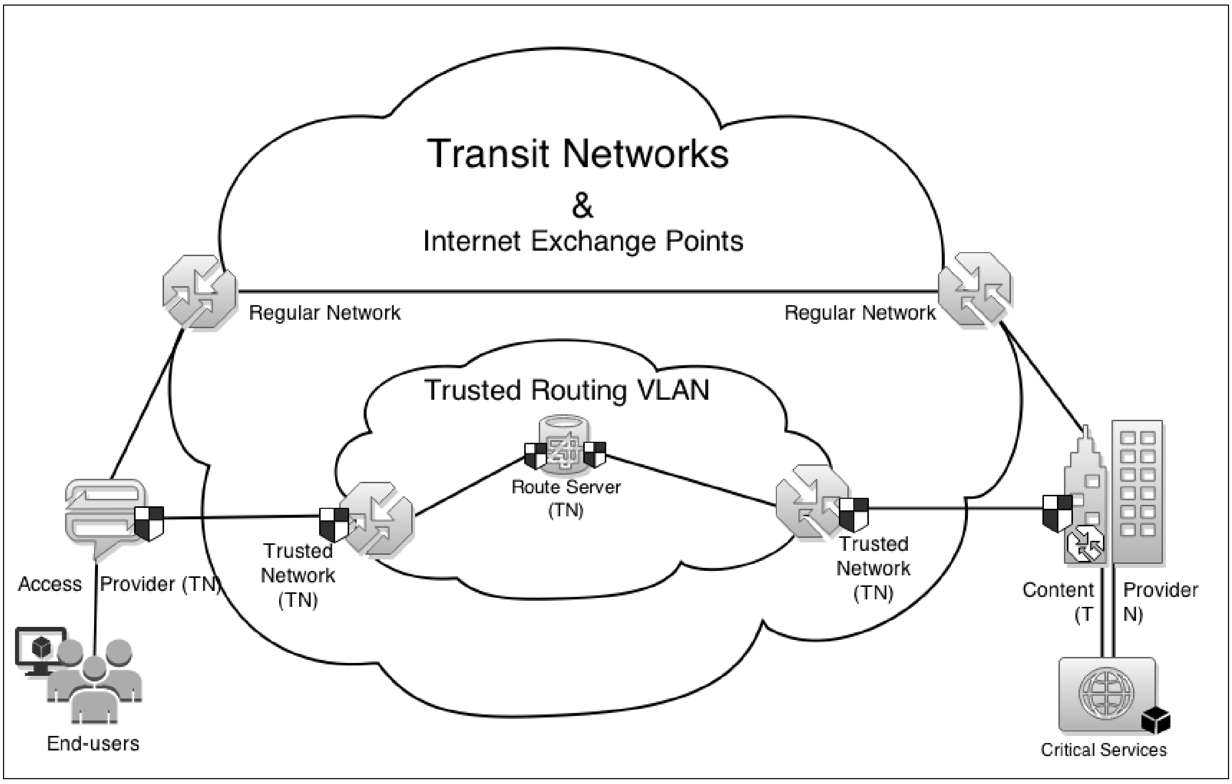

Figure 5.4: TNI Overview

with that party through Trusted Routing and the remaining global traffic will never reach the target.

The second scenario is an always-on solution, which allows parties to route traffic through the Trusted Routing VLAN if the destination party is also part of the Trusted Routing group. If the destination party is not part of the Trusted Routing group, traffic will be routed through the Transit in its traditional form. In the case of a DDoS attack, the participant can choose to disconnect from the Transit connection used to access the internet. Therefore, remaining reachable to the participants of TNI who want to access their services, critical or not. This implementation method is proposed as ideal for governmental, financial and energy critical services since these are mainly aimed towards their national citizens. This scenario could also be automated. An automated implementation could use the maximum load parameters of the router(s) and firewall(s) in order to define an acceptable traffic load threshold. As a result end-users should not notice any difference since the traffic is preferably routed over the Trusted Routing VLAN.

The Trusted Networks Initiative is designed as an extra tool to help retain the availability of the participants in the case of a future attack, which could be either too large to mitigate, last too long, or even both. However, a few questions arise from the aforementioned scenarios. Since the Trusted Networks Initiative should be an emergency measure, what would be the acceptable definition of an emergency for the participants to activate Trusted Routing? Also, by activating only in the case of an emergency, how would the participants regularly test proper working of the emergency-connectivity? These questions have yet to be answered in the Policies and by the members of the initiative.

In addition, this initiative is only realistic if all prominent Access and Content Providers participate as a whole. At the very least, a large majority of end-users of Access Providers should still be able to reach critical services in case of a mid-sized DDoS attack and not be excluded from accessing critical services.

5.5 Routing Manifesto and Project FENIX

The Trusted Networks Initiative is not the first collaboration of its kind which aims to assist in the fight against DDoS attacks. Project FENIX, from the Czech Republic is a similar initiative which is already implemented. And the Routing Manifesto, from the Internet Society (ISOC), is a set of routing rules, meant to be used as guidelines

Project FENIX

Project FENIX is created by the Content Provider NIX.CZ in Czech Republic. The purpose of the initiative is to ensure the uninterrupted operation of Internet services for connected entities during DDoS attacks[29]. The members of the initiative manage themselves independently. This initiative was designed for companies that provide important services and to guarantee connectivity to these entities in order to secure their operation even in the most critical situations [30]. Although Project FENIX is fairly similar to the Trusted Networks Initiative, it has stricter requirements. In order to participate, the organization must sign their domain(s) with DNSSEC. Furthermore, BGP sessions must be protected against session hijacking. Participants must also be part of a Remotely-Triggered Black Hole (RTBH) filtering system[45]. This system allows a participant that has become the target of an attack to identify and block the traffic[29]. Such system is not defined in the Trusted Networks Initiative.

Routing Manifesto

Mutually Agreed Norms for Routing Security (MANRS) can be seen as the fundamentals of the Trusted Networks Initiative[28]. MANRS describes a set of recommendations that – according to their documentation – should definitely be implemented by Content and Access Providers. The Routing Manifesto should prevent the propagation of incorrect routing information, traffic with spoofed IP addresses and facilitate global communication and coordination between network operators[28]. The Trusted Networks Initiative takes this a step further by defining Trusted Networks and implementing an overlay VLAN within the existing infrastructure.
Furthermore, the Trusted Network Initiative allies strongly with the MANRS project by listing MANRS members at the initiative's member-site, and by adding 'awareness of Routing Manifesto norms' as an element of the Trusted Routing Policy.

Chapter 6

Recommendations

During the course of this research, the focus was on gathering information about the initiative in order to give it shape and a provide a better understanding. As previously stated, the research was conducted by communicating with the participants and analyzing their roles, studying the produced Proof-of-Concept and designing separate scenarios, in which TNI could function and theorize about its behaviour and its consequences on service availability to end-users. The intentions of TNI were investigated and for further implementation of the Trusted Networks Initiative, certain remarks need to be addressed in order to improve the service it is offering.

6.1 Policies

First and foremost, the policies[31] of the Trusted Networks Initiative, on which the participants need to agree, are still in a beta phase. The members of TNI remark that the policies are the core and basis of their successful cooperation, therefore, the policy documentation possesses room for further improvement:

- The Trusted Networks Initiative Policies do not define a specific time-frame to defend against a DDoS attack from within the Trusted Networks. Therefore, a time-frame of a maximum of one hour is recommended for the appropriate action to be taken, before implications to the originating party's membership occur.

- Policy 2.4.5 states the need for monitoring traffic in terms of flows and packets in order to detect and signal irregularities in the monitored values. However, it is not clear which level of monitoring is meant. This policy could suggest deep-packet inspection while – according to the members – this is not the case. Nevertheless, to avoid legality issues, vague terms should be avoided in order to ensure the peaceful cooperation of the members.

- Policy 4.1 states audits are carried out in order to ensure network compliance. However, it is not yet clear who the authority conducting said audits will be, nor how they will be financed. Furthermore, not all participants would comply to a third party auditing their network operation and activity.

- Policy 8 states an anti-trust policy. Lawsuit threats do not favour the trust of this initiative. As the participants are basing the success of this endeavour in relationships of trust, potential lawsuits would not favour the trust of this initiative. Due to this fact, a closer inspection of this article should benefit the future of these relationships.

- Finally, it is recommended to revisit the policies once the Trusted Networks Initiative is operative, since, at that point, the participants would have a much firmer grasp on how a large scale collaboration between differently oriented parties can succeed.

6.2 Technical Aspect

As the technical aspect of the Trusted Networks Initiative could not be physically researched and tested, remarks can be derived only by the study of its Proof-of-Concept and the information gathered by members of the initiative, who feel quite confident about the technical implementation of TNI.

Since the Trusted Networks Initiative is similar to Project FENIX and has based some of its ideas on FENIX's manner of operation, a possibly beneficial, for the entirety of the members, suggestion would be for the participants to sign their websites with DNSSEC[37]. Furthermore, BGP sessions should be protected against session hijacking through TCP signatures (RFC2385)[38].

A joint venture between the Trusted Networks Initiative and the non-profit initiative NaWas (Chapter 4.3) is theoretically possible and should be studied. This way, members of the initiative could use a traffic washing solution, even at the unlikely event of an attack coming from within. Furthermore, it would provide the participants with the option of having an alternative to the commercial mitigation solutions.

Moreover, as businesses and governmental departments are centralized around the digital platform, it is essential that their digital platform is always accessible. To realize this, it is important for the critical services, protected by this initiative, to be accessible, not only through the Access and Content Providers' Trusted Networks, but also from Mobile Carriers. It would be beneficial for both participants and end-users to witness Mobile Carriers joining the Trusted Networks Initiative, since more and more end-users are relying on mobile networks and connection when accessing critical services. In addition, the fact that critical services make use of mobile networks for security reasons during transactions and authentication procedures, such as SMS messages, renders the need for their participation all the more important.

6.3 General

Lastly, the members need to reach a consensus on the form of this initiative. Its intention is to be a temporary and last resort solution, however with permanent activation, its effectiveness may be assured on a greater scale than with the on-emergency activation. This is a debatable topic, which may affect the way the members see this initiative and whether or not they will support it, if a major change on the initiative's concept is made in the future. Due to this reason, its members are recommended to agree shortly on the Trusted Networks Initiative's final form before further implementation.

Chapter 7

Conclusions

Throughout the previous chapters of this report, the approach to this research was explained. Background information on DDoS attacks was provided, while the causes and severity of these attacks were stated. The different types of DDoS attacks were described and the current methods of preventing and mitigating them were examined, while also mentioning their benefits and drawbacks. Afterwards, the topic of the research was introduced and the concept of the Trusted Networks Initiative was explained. Its policies – as this is where the members base the success of their collaboration – were studied and the inner workings of the initiative were described. Two different scenarios of operation were given as examples and the relation of TNI to other similar projects was mentioned. Moreover, several recommendations for the future benefit of the initiative were provided and derived through the course of the research.

The research mainly posed the question of the feasibility of such a concept in protecting its members from large and long lasting DDoS attacks. To address this, several sub-questions were designed as seen in Chapter 2, which examined the initiative from a Policy, Technical and End-User aspect and were answered in the remaining chapters.

Unfortunately, DDoS attacks cannot be left in the past. Such attacks are unavoidable and threatening with constant increases in their bandwidth. Nowadays, modern society should be prepared for a future attack which will be beyond the current means of mitigation. Adding resources in terms of bandwidth is not a viable long term solution. The only way to achieve progress in this subject, is by having a variety of tools which can be used when deemed necessary. The Trusted Networks Initiative intends to provide the aforementioned variety. As a temporary and last resort solution, it aims to form a security building block with the currently existing mitigation solutions. Its participants can expect to use their already existent infrastructure – perhaps with a few minor changes in equipment – in order to connect to the Trusting Routing VLAN and communicate with the other Trusted Networks.

Ultimately, the desired goal is to maintain a high degree of availability even under the worst circumstances. To succeed in its purpose, the Trusted Networks Initiative relies on the support of the Internet Exchanges and the large Access and Content Providers. This initiative requires the participation of these members to reach the public and be effective in its task. The goal of the current mitigation solutions is the same, however at times, it can prove to be quite costly[22] even for large organizations. Moreover, through these solutions, the attackers are almost never identified. By enforcing ingress filtering techniques, such as BCP 38, the Trusted Networks Initiative can combat location anonymity within the networks of its members. This has as a result the identification of a possible source of attack within a Trusted Network, which could potentially lead to a criminal court case against the attacker.

Challenges can arise when a web service is deployed in a distributed or cloud-based model. Due to this reason, organizations should align their web services with Trusted Content Providers or either host in-house. In case of a severe DDoS attack, Transit connectivity is temporarily disabled. Therefore, only the connection to other Trusted Networks is possible. The organization should ensure their IT services are hosted within the range of the Trusted Networks.

The key point is that the members of TNI want their services to remain available to the end-users and have their function uninterrupted by digital disruption attacks. The Trusted Networks Initiative could, at least, ensure connectivity between its members, therefore having their services available to the end-users connected through TNI. The above, leads us to the conclusion that the Trusted Networks Initiative could be, in fact, a feasible additional solution against large and long lasting DDoS attacks, however this will depend on the consensus of the initiative's final form. The following chapter suggests future research which could be conducted on this subject.

Chapter 8

Future Work

In conclusion, the Trusted Networks Initiative is a new feasible concept which is interesting enough to be researched further from an academic point of view. It would be particularly thought-provoking to investigate the details of the actual functioning by further Proofs of Concept. Therefore, a research with a test network and/or network simulations could evaluate more specifics in the effectiveness of this initiative. As we did not have the opportunity to be facilitated in an environment, in which we would be able to test the inner workings and operations of the initiative and its PoC, it would be particularly interesting to investigate its successful operation with proper experimentation. A global deployment should reveal how end-users experience the availability of critical services versus no end-users being able to reach the critical service at all.

Furthermore, it would also be interesting to test the effects of this type of routing on services such as e-mail and DNS servers; or to analyze the possible scalability issues which may occur with the BGP switching mechanism.

In addition, a potential effect on net neutrality concerning the end-users could be researched, especially when members will choose to isolate themselves from the outside traffic, making discrimination accusations a possibility. Of course the purpose of such research would not be to simply avoid said accusations, rather to affirm that there will not be a basis for these accusations to be made. The initiators and members of the Trusted Networks Initiative allege of no such intention, since any member solely on its own decides when to step back to 'Trusted Routing only', which is an alternative for not being reachable at all. Thus, the intention is not to restrict certain users from accessing critical services, rather than to ensure the fact that as many users as possible will still be able to access critical services. Nevertheless, ethical concerns will always exist. Therefore, the goal of this initiative should not become its means.

As the Trusted Networks Initiative is reaching its implementation phase, it remains quite a broad subject of research. Both from a technical perspective and ethically point of view, it is a topic worthy of further investigation. We hope our and possible future work will help improve this initiative in its ultimate goal and in the fight against DDoS attacks.

References

[1] Prolexic Security Engineering and Research Team Quarterly Global DDoS Attack Report, 2014, http://www.prolexic.com/kcresources/attack-report/attack_report_q214/Prolexic-Q22014-Global-Attack-Report-A4.pdf

[2] National Cyber Security Strategy (NCSS), From awareness to capability, 2014, https://english.nctv.nl/Images/national-cyber-security-strategy-2_tcm92-520278.pdf

[3] Cisco Guard DDoS Mitigation Appliances - Defeating DDOS Attacks, 2004, http://www.cisco.com/c/en/us/products/collateral/security/traffic-anomaly-detector-xt-5600a/prod_white_paper0900aecd8011e927.pdf

[4] Internet Denial of Service: Attack and Defense Mechanisms, Chapter 3 History of DoS and DDoS http://users.atw.hu/denialofservice/ch03lev1sec3.html

[5] S. Hoffman, DDoS: A Brief History, March 25, 2013, http://blog.fortinet.com/post/ddos-a-brief-history

[6] National Cyber Security Strategy (NCSS), Cyber Security Assessment Netherlands, CSAN-4, 2014 https://www.ncsc.nl/binaries/content/documents/ncsc-en/current-topics/news/cyber-securty-assessment-netherlands-4-cybercrime-and-digital-espionage-remain-the-biggest-threat/1/CSAN%2B4.pdf

[7] Arbor Networks Worldwide Infrastructure Security Report, Volume IX, 2014, http://pages.arbornetworks.com/rs/arbor/images/WISR2014.pdf

[8] Incapsula, The Top 10 DDoS Attack Trends, 2014, https://www.imperva.com/docs/DS_Incapsula_The_Top_10_DDoS_Attack_Trends_ebook.pdf

[9] A. Robachevsky and B. Overeinder, Ripe, Seven Years of Anti-Spoofing, https://ripe66.ripe.net/presentations/168-RIPE_66_Anti-Spoofing_Panel.pptx

[10] cmand.org, Spoofer Project: State of IP Spoofing, 2005, http://spoofer.cmand.org/summary.php

[11] IBM Security Systems, IBM X-Force Threat Intelligence Quarterly, 2014, http://public.dhe.ibm.com/common/ssi/ecm/wg/en/wgl03062usen/WGL03062USEN.PDF

[12] Verisign, Distributed Denial of Service Trends Report, Q3 2014, https://www.verisigninc.com/assets/report-ddos-trends-Q32014.pdf

[13] Neustar, Annual DDoS Attacks and Impact Report - The Danger Deepens, 2014, http://www.neustar.biz/resources/whitepapers/ddos-protection/2014-annual-ddos-attacks-and-impact-report.pdf

[14] Spamhaus, The Spamhaus Block List http://www.spamhaus.org/sbl/

[15] M. Prince, CloudFlare, The DDoS That Knocked Spamhaus Offline, 2013, https://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho/

[16] NL-ix, Various Severe DDoS Attacks in Holland, April 2013, https://www.nl-ix.net/docs/newsletters/NLixNews20.pdf

[17] Akamai, State of the Internet Report, Q2 2013, http://www.akamai.com/dl/documents/akamai_soti_q213.pdf

[18] Cisco Security Intelligence Operations, A Cisco Guide to Defending Against Distributed Denial of Service Attacks, http://www.cisco.com/web/about/security/intelligence/guide_ddos_defense.html

[19] Verisign, DDoS Protection Services Overview, 2012, http://www.nl-ix.net/docs/verisign/faq-ddos-protection-services.pdf

[20] CloudFlare, Keyless SSL, https://www.cloudflare.com/keyless-ssl

[21] Incapsula, Guide To Selecting A DDoS Solution, 2014https://www.imperva.com/docs/DS_Incapsula_Guide_To_Selecting_A_DDoS_Solution.pdf

[22] Incapsula, DDoS Protection Datasheet, 2014http://www.incapsula.com/datasheets/ddos-protection.pdf

[23] Nationale Beheersorganisatie Internet Providers, Beveiliging tegen DDoS aanvallen http://www.nbip.nl/diensten/nawas-demand-beveiliging-tegen-ddos/

[24] Abuse Information Exchange, AbuseHUB van start: botnets aangepakt, November 2013, http://www.abuseinformationexchange.nl/mm_uploads/AbuseHUB_van_start_botnets_aangepakt-1.pdf

[25] Verisign, DDoS Protection Services Overview, 2014, https://www.verisigninc.com/assets/datasheet-ddos-overview.pdf

[26] M. Prince, CloudFlare, The DDoS That Almost Broke the Internet, March 2013, https://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/

[27] RFC 2827, Network Ingress Filtering - IP Source Address Spoofing, 2000 https://tools.ietf.org/html/bcp38

[28] Routing Resilience Manifesto (MANRS), 2014, http://wp.internetsociety.org/routingmanifesto/wp-content/uploads/sites/14/2014/09/MANRS-PDF.pdf

[29] NIX.cz, The FENIX Project Connection Rules, http://nix.cz/cs/file/RULES_FENIX

[30] NIX.cz, The FENIX Project, http://fe.nix.cz/en/

[31] The Hague Security Delta, Trusted Networks Initiative, Trusted Networks Policy, https://www.thehaguesecuritydelta.com/images/20141124_Trusted_Networks_Policy_beta-vs0_7.pdf

[32] The Hague Security Delta, Trusted Networks Initiative, FAQ Trusted Networks Initiative, https://www.thehaguesecuritydelta.com/images/FAQ_TNI.pdf

[33] The Hague Security Delta, About HSD, https://www.thehaguesecuritydelta.com/about-hsd

[34] Open Peering, BGP4 Network - Router Hardware, http://www.openpeering.nl/routers.shtml

[35] RFC 6192, Protecting the Router Control Plane, 2011 https://tools.ietf.org/html/rfc6192

[36] Trusted Networks Initiative, ASP4ALL, Proof of Concept Trusted Routing, July 2014.

[37] O. Kolkman, Nlnet Labs, DNSSEC Howto, 2009, http://www.nlnetlabs.nl/projects/dnssec/

[38] RFC 6192, Protection of BGP Sessions via the TCP MD5 Signature Option, 1998, https://tools.ietf.org/html/rfc2385

[39] Watchguard, Anatomy of a DDoS Amplification Attack, David Piscitello, ICANN SSAC Fellow https://www.watchguard.com/infocenter/editorial/41649.asp

[40] Distributed Denial of Service: Taxonomies of Attacks, Tools and Countermeasures, Stephen M. Specht, Ruby B. Lee, Princeton University, 2004 http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.133.4566&rep=rep1&type=pdf

[41] A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks, Saman Taghavi Zargar, Member, IEEE, James Joshi, Member, IEEE, and David Tipper, Senior Member, IEEE, 2014 http://d-scholarship.pitt.edu/19225/1/FinalVersion.pdf

[42] NSFOCUS, Common DDoS Attacks http://www.nsfocus.com/uploadfile/Product/ADS/DDoS%20FAQ/Common%20DDoS%20Attacks.pdf

[43] Huawei, Botnets and DDoS Attacks Report, 2013 http://www.google.nl/url?sa=t&rct=j&q=&esrc=s&source=web&cd=2&cad=rja&uact=8&ved=0CCsQFjAB&url=http%3A%2F%2Fenterprise.huawei.com%2Filink%2Fenenterprise%2Fdownload%2FHW_315881&ei=3ZApVfGOOJSv7AagvIGAAw&usg=AFQjCNHb_C-DKSdne539rV-mi99CD8ILsQ&sig2=5Iv_Pdcn-P2pvK4q9OUY0w&bvm=bv.90491159,d.ZGU

[44] Sprint ATL Report, DDoS Mitigation via Regional Cleaning Centers, Sharad Agarwal, Christos Tryfonas, Travis Dawson, 2004 https://ipmon.sprintlabs.com/publications/uploads/RR04-ATL-013177.pdf

[45] Remotely Triggereed Black Hole Filtering - Destination Based and Source Based, Cisco White Paper, 2006 http://www.cisco.com/web/about/security/intelligence/blackhole.pdf