
UNIVERSITY OF AMSTERDAM

UNIVERSITY OF AMSTERDAM
MASTER SYSTEM AND NETWORK ENGINEERING
RESEARCH PROJECT

Preventing Common Attacks on Critical Infrastructure

Students:
Wouter MILTENBURG
Wouter.Miltenburg@os3.nl

Koen VEELENTURF
Koen.Veelenturf@os3.nl

Supervisors:
Jaya BALOO
Jaya.Baloo@kpn.com

Oscar KOEROO
Oscar.Koeroo@kpn.com

February 8, 2015

Abstract

Critical infrastructures are attacked just like any other network, but attacks are getting more sophisticated and can not be solved on its own. When outages or malfunctions can have a devastating impact on a regional, national, or international level, an infrastructure can be classified as a critical infrastructure. By interviewing a couple of companies we identified what the most common attacks are that the critical infrastructures suffer from, and why some common practices are implemented or are not implemented. A guideline is provided for network administrators of critical infrastructures on how they can mitigate common attacks on their network.

Acknowledgements

We would like to express our special thanks of gratitude to our supervisors Oscar Koeroo and Jaya Baloo of KPN, who gave us the opportunity to do this project. They helped us a lot during our research by making time for us and helping us getting in contact with third-parties. Secondly, we would like to thank the people we interviewed, namely Rob Vercouteren, Jeroen Veen, Michel Zoetebier, Erik Bais, and Marc Gauw. These people gave us a picture of what kind of problems there are on the Internet. Finally, we would like to thank Arno Bakker that helped us by providing us with helpful feedback.

Note for reader:
Certain commercial entities or equipment are named in this research paper in order to provide configuration examples. However, this does not imply that this is the best available material or equipment for this purpose.

Contents

1 Introduction **4**
1.1 Related work . . . 4

2 Research questions **5**

3 Approach **6**

4 Data gathering **7**
4.1 KPN . . . 7
4.1.1 Jaya Baloo and Oscar Koeroo . . . 7
4.1.2 Rob Vercouteren . . . 8
4.1.3 Jeroen Veen . . . 10
4.1.4 Michel Zoetebier . . . 11
4.2 A2B Internet . . . 12
4.3 NLnet . . . 13
4.4 Multinational company . . . 14
4.5 Identified problems . . . 16

5 BGP Hijacking Analysis **17**
5.1 Peer policies . . . 18
5.2 BCP 38/84 for BGP . . . 18
5.3 Maximum routes accepted . . . 19
5.4 Max prefix length . . . 19
5.5 IRR . . . 20
5.6 AS_PATH filtering . . . 21
5.7 Securing the BGP session . . . 21
5.8 BGP Origin Validation . . . 22
5.9 BGPsec . . . 23

6 DDoS Analysis **25**
6.1 Scrubbing . . . 25
6.2 BCP 38/84 for DDoS . . . 26
6.3 Intrusion Detection Systems . . . 26
6.4 NetFlow . . . 27
6.5 Unicast Reverse Path Forwarding . . . 27
6.6 BGP FlowSpec . . . 28
6.7 Trusted Networks Initiative . . . 28

7 Phishing Analysis **31**
7.1 Sender Policy Framework . . . 31
7.2 DomainKeys Identified Mail Signatures . . . 33
7.3 Domain-based Message Authentication, Reporting & Conformance . . . 34

8 "No business case" **36**
8.1 Return on security investment . . . 36
8.2 Law . . . 36
8.3 Reputation . . . 37
8.4 Awareness . . . 37

9 Additional security measures **38**
9.1 CERT . . . 38
9.2 Responsible disclosure . . . 38
9.3 Monitoring . . . 39
9.4 Patching . . . 39
9.5 Asset management . . . 39

10 Conclusions **41**

11 Future work **43**

12 References **44**

Appendices **48**
Appendix A Network Diagram . . . i
Appendix B Configuration example route filtering . . . ii
Appendix C Configuration example maximum prefixes . . . iii
Appendix D Configuration example AS_PATH filtering . . . iv
Appendix E Configuration example MD5 and IPsec . . . vii
Appendix F Configuration example BGP Origin Validation . . . x
Appendix G Configuration example uRPF . . . xii
Appendix H Configuration example SPF . . . xiii
Appendix I Configuration example DKIM . . . xiv
Appendix J Configuration example source filtering . . . xv

Chapter 1

Introduction

Right now, somewhere on the Internet a network is being attacked. Some are easier to mitigate than others. However, when such an attack is successful, it can have a devastating effect for other people. This depends on the type of attack and what kind of services the company provides.

It is necessary for companies to have the necessary security mechanisms in place to successfully mitigate an attack, since they are getting more advanced. This research report focusses on the attacks against critical infrastructures, which is in this case defined as an infrastructure where outages or malfunctions can have an impact on a regional, national, or international level.

There are already a lot of tools or common practices to make a network infrastructure more robust to the most common attacks. Therefore, it is interesting to see why these mechanisms are not implemented. To get a better understanding we also investigate the steps that need to be taken before a company is willing to implement a certain measure. This will be done from a technical perspective and from a business perspective as well, since implementing new techniques takes time and will cost money.

Gathering information about the mechanisms that are implemented by companies, which serve critical infrastructure, is done by interviewing them. Information that has been gathered during these interviews is summarised and made available for the interested reader in this report. Companies will also be asked why certain techniques have not been implemented and what needs to be done to make the Internet, and therefore critical infrastructure more robust.

The research is focussed on identifying the most common attacks on critical infrastructure, what can be done to harness such a company, and see why certain mechanisms are or are not implemented. The report will also be a guideline for network administrators on how to implement the suggested security measures.

It is expected that the interested reader is familiar with the concepts of BGP and basic understanding of the mail system.

1.1 Related work

A lot of work has already been done for protecting a network for certain attacks on the Internet, targeted at critical infrastructure. One of the attacks that networks in general suffer from are DDoS attacks. A possible solution could be the Trusted Networks Initiative [61] that aims at providing a last-resort measure when a network is suffering from a DDoS attack.

One other common attacks for networks in general is BGP hijacking. Work has already been done to prevent this and several methods exist. The two methods that the Secure Inter-Domain Routing (SIDR) working group is working on is Origin Validation and BGPsec. BGP Origin Validation has been standardised by the IETF in RFC 6811 [38] by J. Scudder et al. and BGPsec is still a draft [28] by M. Lepinski.

However, more common attacks exist and would be gathered trough interviewing. When related work has already been done it would be referenced when the work is outlined.

Chapter 2

Research questions

The Internet as we know it today is not safe by design. The protocols used today are not designed with security in mind. KPN wants to get a clear picture of how attacks can be prevented and how the Trusted Networks Initiative could help to mitigate DDoS attacks. The Trusted Networks Initiative initiated by The Hague Security Delta and NLnet tries to create a measure that can be used during emergencies. Instead of creating temporary solutions, we should tackle the problem by the source, and improve the overall security of the current Internet infrastructure, which is not only limited to DDoS attacks. Since DDoS attacks are only a subset of the Internet problems, it needs to be investigated what kind of other attacks are out there.

The research is formed around the following main research question.

Which techniques are available today that could be used to mitigate common attacks?

As result of the main question, shown above, the following sub questions were formulated.

- What kind of problems does the Trusted Networks Initiative try to solve?
- What kind of attacks are critical infrastructures suffering from?
- What kind of techniques can be used to harness yourself or your company (e.g. ingress/egress filtering, trusted routing, etc.) to protect against certain attacks that the critical infrastructures suffer from?
- If there are techniques available, why are these not used in common practice?
  - Which step(s) could assist adopting these techniques?

Chapter 3

Approach

During the research the security of critical infrastructures was investigated. There are different techniques that can be used to mitigate common attacks. Common attacks are defined as attacks that most companies suffer from. The companies KPN, A2B Internet, NLnet, and a Dutch multinational company were interviewed to get a picture of what is happening on their network according to them, what kind of security measures they implement and how these companies implement them. A list of common attacks forms of abuse, and the measures that companies take was created, based on the information gathered by interviewing the companies.

During the theoretical research part of the project, the information of the interviews was used to look for available techniques that can be used to mitigate attacks. If techniques found during the theoretical research were not used by the interviewed companies, the reason or reasons for not adopting these techniques were researched, as well as the steps companies should take to adopt these techniques.

Chapter 4

Data gathering

A couple of companies were interviewed for this project. The supervisors of this project from KPN helped setting up interviews with several people from KPN, including people from the Chief Information Security Office (CISO) department and the Computer Emergency Response Team of KPN (KPN-CERT). We also talked with other parties like A2B Internet, NLnet, and a Dutch multinational company. The CISO department of KPN is responsible for the internal security policies, penetration testing of hard- and software, and gives advice within the company on several security related topics. KPN-CERT is part of the CISO department for incident handling.

All of the information outlined in the following sections, is data collected during the interviews. Possible solutions that are discussed in these section are proposed by the interviewees. In the next chapter we will propose solutions for the identified common attacks.

4.1 KPN

A couple of interviews were held with persons of one of the largest Internet Service Providers of The Netherlands [4], KPN. KPN offers (mobile) Internet, landline and mobile phone connections, and television for private customers and businesses. A lot of companies are dependent on Internet connectivity by KPN (e.g. banks, insurance companies, and governmental services). These companies will suffer great loss when Internet connections are disrupted. Therefore, KPN maintains one of the critical infrastructures in The Netherlands, because KPN serves a majority of the commercial home Internet connections, and other important infrastructures, e.g. the national emergency service 112.

4.1.1 Jaya Baloo and Oscar Koeroo

The initial interview for this research project was with the CISO of KPN, Jaya Baloo and one of the team members Oscar Koeroo [33]. The subject of the interview was mostly about what kind of problems KPN runs into when maintaining their infrastructure and what the main challenges are on the Internet. According to Jaya Baloo it is very hard to make large ISPs and Internet companies implement common practices. For instance, on the subject of IPv6, the Internet world is talking a lot about it for years, and that "we should implement it!". However, most Dutch ISPs did not turn on IPv6 on their network for customers and they are stalling the further implementation of it.

> The fundamental problem on the Internet today is: "How can we trust someone on the Internet?" (From: [33])

Therefore, how can we be sure that I am talking to 'Bob' and not to 'Eve'? This is a problem that we constantly try to solve, and we need to be aware of this security risk when we use routing protocols like BGP.

Not all Internet companies implement techniques like BCP 38. Therefore, it is possible for users to spoof their source address. A result of that, is that DDoS attacks are quite easy to realise. One very important aspect of mitigating DDoS attacks within a network is the response time of the 'digital firefighters'. The time between the starting point of the attack and the detection of an

attack is also the opportunity time for attackers, so how much time does it take for an attacker to infiltrate a network, and how long can the attacker stay undetected. However, DDoS is not the main issue. How can we be sure that we are talking to our neighbours? We simply can not trust each other on the Internet. One thing that helps to increase the trust is to agree on strict routing policies. In these policies peers agree on which routes to announce to each other. KPN adopted the Mutually Agreed Norms for Routing Security (MANRS) of the Routing Resilience Manifesto [47]. The most important actions described in the MANRS are preventing propagation of incorrect routing information, preventing traffic with spoofed source IP addresses, facilitating global operational communication and coordination between network operators, and facilitate the validation of routing information on a global scale.

DDoS attacks are relatively easy to mitigate, using scrubbing services and filtering. Routers used today, are not able to handle fine-grained filtering with hundreds of Gigabits of network traffic volumes, so the only possibility left is rough filtering. The main task of ISPs is the transportation of Internet traffic. However, due to the amount of network traffic generated by DDoS attacks, ISPs need to filter out this traffic. The main responsibility however, lies at the end-users and they should do their best not to participate in a DDoS attack. The end-users' responsibility and how they can harness themselves is not in the scope of this project. Unfortunately, there is no single solution to solve all problems in one go. In this report the focus is on the most common attacks on critical infrastructures.

> There is no such thing as a silver bullet! (From: [33])

One thing that helps a lot in tackling problems is the Delphic maxim "know thyself". "Know your good qualities and your bad ones". According to Jaya Baloo there are over a million IP addresses and millions of hosts used by KPN and its customers. It is almost impossible to know what all IP addresses or hosts are for, what kind of software they are running, which version of that software is running, and what they are doing exactly. In case of KPN, they are trying to actively scan for abusive users in order to take them of the network.

> Nine out of ten times, we can only find the country of origin. (From: [33])

A trend of the last couple of years is the certification of devices, "when a device has a sticker that says everything is okay, it is safe". This kind of certification can be very misleading. If the policy is to have an up-to-date operating system on a device, but the applications on that device are not up-to-date, the device is not per definition secure and safe. There is one very important aspect of security from a management point of view. When the management of a company does not reserve money for security, there will be no security. Unfortunately, security will only cost money. However, security measures could save a company money as well. When a company has a lot of bad publicity, and it will eventually get a bad reputation, it is possible that the company will lose money over a minor-looking issues. Since 'the hack of 2012' of KPN [15], the priority of security in general has been upgraded to a higher level in the company's hierarchy. From 2012 on, the CISO department has more influence on the day-to-day operations.

4.1.2 Rob Vercouteren

One of the first interviewees was Rob Vercouteren, a network specialist of KPN-CERT [36]. He has around seventeen years of experience at KPN with network engineering. The network of KPN is

designed with a couple of layers. KPN maintains a very strict policy on network design and placing systems in different zones. Using this layered design, it is more difficult to access for instance the database server from the external network. One would need to go through each layer before getting to the actual data. This behaviour would likely be detected by an IDS or a monitoring system. Besides the network and system design policy, KPN has a variety of security policies on various topics, including routing policies. These policies help to standardise the architecture and overall platforms within KPN, which improves the overall security.

One of the major problems the Internet suffers from in Rob's opinion is the lack of anti-spoofing measures. Up to this day it is very easy to spoof traffic and send it over the wire. There are a couple of measures that could be taken to prevent the use of spoofing, but a lot of those measures are not implemented by Internet companies. A solution proposed by Rob, would be the use of ingress and egress filtering. Unfortunately, there are companies (e.g. bullet proof hosting companies) that will not implement this, because they are making profit from accepting spoofed traffic over their network. Another reason for not implementing this kind of filtering is the costly process in terms of CPU cycles. A lot of equipment currently used by network operators would not be able to handle the extra load on the system for filtering. Especially with routers that handle high volumes of network traffic, e.g. ISPs like KPN.

> On our network, every attack that one can think of comes by, e.g. DDoS, port scans, and sniper attacks. However, it is likely that there are many other attacks, but that would mean that we have to look into Internet packets. This is only allowed with full permission of the customer. (From: [36])

An example of an incident at KPN from a couple of years ago, was a case where someone hijacked a more specific of an IP superblock. This resulted in over two thousand customers not having an Internet connection any longer. When this happens, one needs to troubleshoot the situation, which requires some knowledge of BGP and the right tools. Therefore, troubleshooting will take a lot of time. A framework for creating trust in the Internet's routing infrastructure is RPKI (Routing Public Key Infrastructure, also known as Resource Certification). The major downside of this framework is the fact that one is dependent on other parties. One can only benefit from this technique when more operators are using Origin Validation using the RPKI. Unfortunately, since there is a lot of debate about the implementation of Origin Validation, especially in the APNIC and ARIN community, there are not that many companies that have implemented Origin Validation. The RIRs APNIC and ARIN force users of RPKI to sign a disclaimer of warranty and indemnification clauses. A lot of people within the community are opposed against signing such an agreement [3].

A problem of BGP is the possibility of hijacking a BGP session, and the possibility of issuing TCP resets. IPsec can be used to make BGP more robust against these attacks. KPN does not use IPsec to protect the BGP session, but uses techniques like MD5 passwords and ACL filtering.

> Unfortunately, there are no real solutions for fixing the Internet problems. The solution still needs to be created. In the ideal situation we should start designing the Internet from scratch. (From: [36])

Nowadays, Internet companies use NetFlow to monitor their network flows. With the use of network flow data, it is for instance possible to track down the entry point of a DDoS attack in one's

network. The newest industry standard is Internet Protocol Flow Information Export (IPFIX) [9]. The IETF IPFIX working group used Cisco's NetFlow v9 as a basis for this new protocol. Flow information is gathered from a router or multiple routers in one's network. According to Rob Vercouteren, NetFlow (or other NetFlow-like solutions like IPFIX or sFlow) is a very powerful and important tool used by Internet companies. He does not know another tool that offers the same feature set as NetFlow, but if there is a company using another tool for monitoring a network he would be very interested, since there are not many other solutions.

For spam issues, KPN uses a reputation system. When a customer sends a lot of spam and gets reported, the customer will get reported in the reputation system. In case of the small business customer, when the reputation of the customer exceeds a limit x, the Internet connection will be null routed. The customer will need to address the spam issue before the Internet connection will be restored.

Besides the abuse on networks, another problem is making mistakes in configurations, e.g. accidentally announcing the wrong routes using BGP [44], and/or accepting large amount of routes, instead of accepting a specific amount of them. These configuration errors can cause a lot of trouble on the Internet.

Another problem is the fact that generally business comes before security. Due to the fact that KPN got hacked in 2012 [15], the board of KPN decided that the cyber security of KPN should have a higher priority. Such decisions are generally based on politics, financial impact, and personal believes.

In The Netherlands, the Dutch government changed the Telecommunicatiewet[1] [2] to include an article about net neutrality. ISPs like KPN are bound by this law to transfer all traffic that is handled by their network. In case of an large attack on the KPN network, KPN is only allowed to block traffic temporarily. When an ISP would block the traffic permanently, users can not reach the source network of the attack. However, there might also be legitimate destinations in that network. When ISPs block traffic on permanent basis, they would break the law and can expect fines by the regulator.

All previous mentioned problems can not be solved on one's own, according to Rob. Therefore, companies need to cooperate with each other, even with their competitors. These collaborations take place through the use of CERTs/CSIRTs (Computer Emergency Response Team/Computer Security Incident Response Team). For instance, KPN has its own KPN-CERT. Together with other CERT teams all over the world, they exchange information about security issues and solutions. According to Rob Vercouteren, this is a very helpful instrument and more (larger) companies should create their own CERT and/or CSIRT. When those companies have their own CERT/CSIRT they are able to react much faster to computer security related issues, and they can learn from each other. Rob pointed out that ENISA (European Network and Information Security Agency) created a guide [17] on how to set up a CSIRT. He recommended that companies that have not created a CERT/CSIRT should take a look at that report.

4.1.3 Jeroen Veen

Since 2001, Jeroen Veen is a network operator and security architect at KPN (formerly Planet Technologies) and has gained a lot of experience in network operations, network design, and security design. In his opinion, one of the biggest problems is that one needs a very good business plan to

[1] The Dutch Telecommunications Act

implement something. If new projects do not have a way of creating profit for the company, it is not likely that the project will launch. However, if new features are requested by a majority of the customers, then it will become interesting for the directors of a company.

> In the case of implementing DNSSEC [5], ISPs do not have a drive to implement it, because customers do not require it from their ISP. In such cases, the government could stimulate the implementation of such techniques by policy. (From: [34])

From his experience at Planet Technologies, the engineers wanted to try new techniques, but in a large company, one needs to make sure the service is stable. It would be a very big issue if, because of implementing new techniques, the original service that should be improved actually suffers in stability and continuity.

In order to detect DDoS attacks, KPN implemented Intrusion Detection Systems (IDSs) within their internal network. They use IDSs to be able to scan the traffic that flows to their systems, e.g. the DNS servers. KPN is not allowed to implement IDSs to scan customer traffic, unless customers (mostly business customers) explicitly ask for it and give permission to scan their traffic.

A wish of Jeroen Veen is to have a central log management system in place. At the moment there is no real central log monitoring at KPN. It would really improve the security of the network, since one can correlate different sources to get a more reliable view of the infrastructure.

4.1.4 Michel Zoetebier

Since 2003, Michel Zoetebier is an abuse specialist at KPN, and is since 2014 Security Officer at KPN-CERT[2]. According to him, spam emails are not a very big problem any longer for ISPs. The spam filters work fairly well. However, phishing is causing a lot of trouble for the normal customers, but also for employees of companies. Through the use of phishing emails, black hat hackers could gain access to computers. If that computer is a company computer, the security of the company network could be in danger. One of the mitigating solutions for this particular problem is the use of SPF, DKIM, and DMARC, according to Michel. With the use of SPF, the mail servers will verify the origin of the email, so spoofed email addresses will be ignored by the mail servers. This will only work if both parties have implemented SPF. For KPN, the fragmentation of mail servers is a problem for implementing SPF and DKIM. All the mail servers need to be located, DNS records need to be created, and additional software needs to be installed or configured.

Another problem that more companies will be suffering from, is that one implements a secure service on an insecure platform. Especially when companies merge, merging the technical parts of the companies will be very difficult and a time consuming process. This results in running multiple (legacy) instances of the same service.

> Your service could be the safest service on earth, if your platform is insecure it could be hacked anyway! (From [34])

Besides the digital security there is also the aspect of physical security. It is very frustrating if one finds an open door to a secure room or is able to walk right into a building without anyone knowing. After the hack of KPN in 2012, KPN changed the hierarchy of the company to make the security department more important. In this way, the CISO department of KPN can advice the management of KPN more directly. In 2013, KPN introduced the CISO REDteam.

[2]KPN's Computer Emergency Response Team

> The [RED]team is involved in security tests of KPN applications and services to ensure that our customers' data is safe from unauthorized access, modification and loss. [...] In addition to the CISO REDteam, KPN is willing to cooperate with anybody who identifies a potential security leak, provided they use their expertise responsibly and respect KPN's Responsible Disclosure code. (From [27])

Since the KPN REDteam, as mentioned above, is also using responsible disclosure, they were able to fix a lot of potential vulnerabilities in their services. According to Michel Zoetebier, since the introduction of the REDteam it has been earning money for KPN instead of losing money. All the possible vulnerabilities the team found, resulted in KPN not having possible reputation damage. When a company suffers from a lot of reputation damage, it might lead to money losses. One could say that the hack of KPN turned out very well for the security of KPN, and its customers. The hack of KPN raised a lot of awareness to the employees of the company and to its customers.

4.2 A2B Internet

We also interviewed Erik Bais from A2B Internet [32]. A2B Internet is a Dutch Network Provider with regional data centres in The Netherlands and Belgium. Erik Bais has years of experience in the Internet world and is an active member of the RIPE community. This interview was mainly to see the other side of the story, since KPN can mainly talk about their ISP network, A2B Internet can talk about the attacks that are targeting their customers and its own network.

One of the most occurring attacks are DDoS attacks where the source addresses are spoofed. Erik Bais categorises this as 'annoying, nothing more and nothing less'. During the interview several methods were discussed to mitigate DDoS attacks. One of the most traditional methods is to use a scrubbing centre that 'scrubs' the traffic and passes it back to its original destination. Filtering a DDoS attack on the edge of a network is not really useful, since the traffic is still coming towards the network and utilises the bandwidth that could be used instead for 'legal' traffic. A more effective method would be preventing malicious network traffic coming to one's network, since it does not utilise the network, but instead it blocks the traffic from the peer it is originating from.

One of the other common attacks is route leakage or rogue route announcements, where an unauthorised peer is announcing a prefix that has not been allocated to him. During the interview it was argued if solutions like Origin Validation and the upcoming BGPsec could solve this problem. It was argued that the biggest problem is the unwillingness of some companies to implement Origin Validation, with BGPsec it is a whole other story since it is still in development, and some additional 'trust' issues that comes along with RPKI [3]. If companies do not implement Origin Validation, or do not use Origin Validation to validate the routes that they receive, there will always be room for these kind of attacks. Such an attack will only affect the ones who have not implemented BGP Origin Validation and do not validate the routes received.

Another feature that Erik Bais pointed out, is the feature that RIPE NCC provides. When someone creates ROAs [57] for its delegated prefixes, the automatic alerting mechanism will alert one when someone else announces the prefixes of that ROA. The alerts can be used for monitoring the allocated prefixes and thus one's network. This aspect also points out that it is really important to have good monitoring in place. Erik told us that it is a good thing that network operators will be alerted when something out of the ordinary happens. During an emergency or an incident, it is important to know what caused the problem, and monitoring can help one to provide the

information that one needs. Without monitoring, the detection and response time would be severely longer.

One of the other problems that Erik Bais pointed out, is that some network operators easily forget the most simple restrictions that they can apply on the routes they receive from their neighbours. One of the examples is that some network operators configure their router to accept a ridiculous high number of prefixes from one of their neighbours. In the case of a neighbour that only sends only tens of prefixes they accept, for example, one thousand prefixes. If this is their only countermeasure that prevents their network against BGP prefix hijacking, the attacker can use this opportunity to announce some more specifics from a compromised network.

Once more, since multiple companies already outlined this, the importance of the business case was stressed. Some businesses make their money out of providing networks that do not implement any kind of spoofing prevention methods. If someone managed to change the company's opinion, so that they implement the anti-spoofing methods, some of their 'customers' will go to another company that did not implement these mechanisms. It was discussed that laws in this case would not really help against source spoofing, since some people do not play by the book.

That is why Erik Bais came up with another solution, called 'naughty ports'. If some peers only send DDoS traffic most of the time, whereas 'legitimate' traffic is only a few megabit per second, one can place them on the 'naughty port'. This naughty port can only handle 100 megabit per second, so when they send DDoS traffic the impact is limited. All legitimate traffic (from the peer where the DDoS attack is originating from) does not notice the 'naughty port' when there is not an ongoing DDoS, but will only be affected when there is a DDoS going on. This is a problem of one's peer, since they did not implement any countermeasures, or good enough countermeasures, to prevent DDoS from traversing their network. It is something that Erik Bais looks into, but from his point of view it looks promising.

4.3 NLnet

To get a better insight in the problems that the Trusted Networks Initiative tries to solve, we had an interview with Marc Gauw from NLnet [35]. NLnet is one of the founders of the Trusted Networks Initiative, hereafter referenced as TNI, and could therefore give more detailed information about TNI and details about the technical implementation.

It was important to get a better overview of the project and Marc Gauw explained what TNI tries to solve exactly. Most companies that participate in this initiative suffered from the DDoS attacks as the Internet suffers from on a regular basis. Attacks can get bigger and bigger, which might result in a DDoS attack that is not possible to mitigate using traditional solutions (e.g. filtering and/or scrubbing as discussed in Section 4.2). This project aims to provide connectivity with a set of peers that can be considered trusted, during such an attack.

As businesses depend more and more on the Internet, it is important that connectivity is still available, even during an attack. Lets consider a bank that needs permanent connectivity with its end-customers for transactions. It is important that the connectivity between the bank and its customers remains during an attack, which is something the TNI aims for. For example, bank A can peer with several ISPs so that customers of the ISPs can be reached using the TNI network. When bank A is being attacked, they can decide to disconnect from the Internet, but still be connected to a significant part of their customers via the ISPs on the TNI network. Transactions between bank A and other banks can still occur as long as they are connected to the TNI network.

One of the common misunderstandings is that companies see this as the one and only solution for solving DDoS attacks. Marc Gauw argues that this is not the case and emphasises once again that it is a last-resort method. Companies should still use the traditional methods, but when there is no other solution available, a company can decide to temporarily disconnect from the Internet.

One of the other common misunderstandings is that TNI is not aiming to provide a solution for other kinds of attacks, besides DDoS attacks. We thought that with the information that was available, it was aiming for more than only providing a last-resort measure for DDoS attacks. During the interview it was explicitly stated that this project is only aiming for providing a last-resort method for DDoS attacks. Marc Gauw argued that it could be possible that the project will be aiming for more, since some protection mechanisms can be implemented in the TNI network, and that it might be implemented in the future. The example that was discussed was to explicitly state in the policy that peers must implement Origin Validation and if that was something that would be considered.

The technical implementation is not that much different from the implementation of peering sessions on an Internet Exchange. Within the TNI network a peer must be connected to VLAN 112[3], and can then connect to the route server. This route server provides the routes of the other peers in the TNI network. The idea is that one receives the routes from all the peers, not that there is some filtering in place to specifically create some routes destined for a certain peer. However, during the interview it was stated that this is possible but it is not the intention of the project.

With some companies, and mostly larger companies, a question can be raised about the peers one should set up a session with. A peer is per definition something equal, within the TNI network everybody is a peer and everybody is per definition equal. This can raise some questions with for example an ISP that would normally provide services to a customer for Internet connectivity, for which the client should pay, but within the TNI network the traffic exchanges free of charge. In the TNI network this is somewhat different, since it is an last-resort method and would only be used on an incidental basis.

Therefore, the focus of this project is more focused on banking, insurance, tax, governmental organisations, and other companies that serve critical infrastructure. Currently, the Trusted Network Initiative concept is in a trial phase. They are currently working on establishing a foundation for TNI, and the next phase involves attracting more national, and maybe international participants.

4.4 Multinational company

One other company that we have interviewed was a multinational company that gave their insights of attacks that the critical infrastructure is suffering from. During the interview we spoke with a member of the security team, from now on referred to as "Mark", of the multinational company about the attacks that they experience and how they harness themselves [37].

Their network is divided in several layers, one of which is the external facing part of the company. This externally facing part, which consists mostly of web servers and mail servers, suffer from most of the attacks. One of the biggest attacks that they are suffering from is spear phishing. Mark told that to protect the company against spear phishing it is mostly important to create awareness in the company itself. These kind of attacks are getting more sophisticated and users should be aware of the consequences if such an attack is successfully exploited. Therefore, the company hires a third-party to send phishing mails to the multinational two times a year. When a person opens the

[3]A pun to the telephone number 112, which is the Dutch/European emergency telephone number.

mail, and clicks on the link, or executes the payload, they will see a video of what just happened. The multinational company also creates awareness inside the company itself. Computers that are accessible in the hallways will show some useful tips and tricks to protect a user. One of them to alert users of the spear phishing mails.

An Advanced Persistent Threat (APT) that the multinational experiences is spear phishing mails. SPF and DKIM can help one in the case of spear phishing mails, since only an authorised mail server can send such an email if validation of the mail is properly done. However, these security measures are not used by the multinational company. Mark told us that most of their IT is outsourced and that there is no priority at the moment to implement SPF and DKIM for their mail servers. Since the users are alerted, and appropriate mail filtering has been installed, the risk is acceptable at the moment. The other aspect of SPF and DKIM is that someone can not send a mail any longer on behalf of one's company. This extra security benefit is not really necessary for the multinational as Mark stated, since their company does not really send a lot of crucial mails to their customers like an ISP does.

Spam on the other hand, is not really a problem for the multinational. The mail servers that they use are outsourced to Microsoft, which have appropriate filters for spam. Most spam mails are filtered out, and executables that have been added as an attachment, will be blocked and can therefore not be executed. Therefore, spam is in itself a low risk where spear phishing is a higher risk.

Due to the layered approach of their network, it is harder for an attacker to get access to the highly confidential data of the company. According to Mark, an attacker has to get through several firewalls and steppingstones to get access to the more classified zones. The external facing part of the company has, most of the time, only port 80 or 443 open for the Internet. However, these machines may still be attacked. When a machine is compromised at some point, they can not reach the confidential zone through the several zones, firewalls, and steppingstones without alerting the IDS.

One of the other attacks that the multinational experiences are the DDoS attacks. According to Mark, the network has been outsourced to a big international ISP, and they can mitigate the DDoS attack early on in their network. However, some parts of their network are connected to other peers as well. In the case if one of the peers causes a DDoS attack, the traffic can be redirected to their own scrubbing centre. This scrubbing centre will deliver them the clean traffic, which the multinational can handle.

Since large parts of the network, for their data centres and uplinks, have been outsourced to the same large international ISP, Mark did not think BGP hijacking was a high risk for their company. If such an attack would occur, it would be something for the ISP to fix. In the case of a more sophisticated attack, where the traffic is redirected to an attacker and then delivered to the multinational, the attacker can not really do much with it according to Mark. All traffic that would be exchanged between the multinational and a client is encrypted. In this case an attacker can only cause harm when the connection to the multinational would be disrupted, which would be something for their ISP to fix.

So far, we have only discussed the inbound DDoS attacks. To protect the network for not initiating a DDoS, by means of an amplification attack by using spoofed source addresses, it is in this specific network filtered out by the firewall. As far as Mark could tell, they do not apply ingress or egress filtering, or uRPF in their office network. However, all DNS traffic has to go through an internal DNS resolver that would respond to the requests. If a client hits a certain threshold, which

it would hit during a DDoS attack in most cases, the SOC team would get an alert. If an attacker would avoid the internal DNS resolver, the firewall would drop outgoing DNS requests and flag it, which will raise an alert. Other types of amplification attacks, like an NTP amplification attack, are treated the same (i.e. blocked by the firewall).

To parse all the logs and create alerts for the SOC team, they would use the Security Information and Event Management (SIEM) system ArcSight. This is a great utility, according to Mark, that really helps the company in identifying the events that really should get their attention. It helped in improving the overall security of the company and it still does. However, to administer such a tool, and perform analysis on the data, it requires time and experience. It is not something that one could install in a few days, it takes time. The risk of such a system is to get overwhelmed with information, which may well be false positives. Therefore, the system needs to be carefully configured, in order to reduce the amount of false positives.

Mark also pointed out that it really is important to share information and collaborate with their competitors. It is a 'give and take' relationship, where people should help other people, and if they do, they will get help back when needed. Sharing information really helped the company, and Mark really advised to set up a CERT/CSIRT and share information with other CERTs.

What currently is on the agenda of the multinational is to deploy NetFlow on as much network equipment as possible. Using this data, and passing it through ArcSight, gives the multinational a better insight into their network. It is something that they have been experiencing for some time, but now made it to their agenda.

The multinational company did not have any major incidents, but did not implement all security measures that they could have done. Security measures that are not currently implemented are DKIM, SPF, DMARC, BGP Origin Validation, and BCP 38/84. Most of this has to do with the business case. At the moment it does not have a high priority, since securing their assets is their most important objective. They would not really benefit from implementing these security measures, it would only cost money and right now the security measures would not win over the business case.

4.5 Identified problems

The next chapter will discuss the identified common attacks and outline how these can be prevented. Together with KPN we identified the following common attacks: BGP hijacking, DDoS, and phishing. However, the next chapter will also discuss some identified problems that do not really have a technical aspect. Something what has been outlined a number of times during the interviews was the importancy of having a good business case for implementing a security measure. When there is no business case at all, it is very hard to defend one's choice of implementing a security measure that costs money.

Chapter 5

BGP Hijacking Analysis

BGP hijacking is an attack where someone announces a prefix that has not been allocated to them. This can occur intentionally or unintentionally, meaning that it is a deliberate attempt to redirect the traffic to their network or that it can happen due to a configuration fault. The attacks can vary from being "annoying" to a more stealthy attack with malicious intentions.

Before elaborating how the attack works and how one can prevent it, it is first necessary to get a basic understanding of BGP. BGP, known as the Border Gateway Protocol, is a routing protocol. This protocol is used by routers to propagate routing information to other routers. It allows networks to update routes whenever the route itself is changed and releases the network operator from creating manual routes to reach other networks. The current BGP version is version 4, which has been standardised in January 2006 [44].

In the era that BGP was originally developed, BGP version 1 is described in [30] and was submitted in June 1989, the developers did not account for the security threats of today's Internet. It is therefore possible that someone can announce a prefix that has not been allocated to that certain entity.

This can not only result in a temporary degraded network availability, where some or all of one's users can not reach the legitimate network and are instead reaching a network that does not belong to the legitimate owner of the prefix, but it can also be used for other attack scenarios.

One of the more advanced attack scenarios could be in the case of temporarily announcing a more specific prefix, possibly a /30, so that traffic destined for these specific IP addresses is diverted to the attacker. It can for example announce this more specific prefix where one of these IP addresses is normally used for a mail server of a company. If the server is configured in such a way that it accepts all incoming mail, it could look at the mails itself and might find some confidential information. If one takes this a step further, the attacker could go to a website, which the attacker knows for sure is used by the company, and could issue a password reset. Most of the time it is only a matter of clicking on a link in the email to reset the password. If the attack time is really short, a couple of minutes, it could be the case that this attack goes undetected.

The attack scenario above truly depends on the configuration between the peers and if they apply some means of filtering. In the following sections we will discuss some of these filtering techniques or other techniques that could be used to detect or prevent this. We also provide configuration examples of the suggested security measures. The network setup where the configuration examples are based upon, is shown in Appendix A.

However, with the suggested security measures in place one's network is protected from accepting rogue routes. All network operators should implement security measures to protect their network from announcing rogue routes or accepting them. If companies will not do this, some networks will still be susceptible for the attacks discussed above. It is therefore important that companies will be stimulated to implement the security measures and be pointed on the risks that they are taking.

5.1 Peer policies

In Section 4.1.2 it is outlined that it is important to have some kind of policy with one's peering neighbours. In this policy it is explicitly stated what the peers expect from each other. This is both related in a technical aspect and a more administrative aspect. First, it is important to formalise the prefixes that will be announced by the neighbour. But it is also important to know how the neighbour can be reached in the case of an emergency. Emergencies can vary from rogue announcements to network outage. It may sound obvious, but it can happen due to job/title changes that someone is not reachable on his old mobile phone number. The policy should also state how information can be updated and what kind of countermeasures the neighbour takes to prevent certain attacks. These countermeasures will be discussed in the next sections and it is important that some basic countermeasures are implemented to prevent source spoofing and advertising rogue routes.

It will take some time to get this formalised policy completely suited for one's company. A good reference to base a peer policy on, is the Mutually Agreed Norms for Routing Security (MANRS) [47] of the Routing Resilience Manifesto. If one is looking for a new peer and has several options, the decision can be made on the basis of a peer policy. Besides the fact that one gets to know its peer, the policy should also state what happens when a neighbour fails to meet the expectations that are stated in the peer policy. Formalising a peering relationship can be a burden when one has a lot of neighbours to choose from, but for a company that wants to carefully decide the peers it wants to connect to, this can be a good tool for making the decision. For network operators with numerous peering relationships it can be quite a task to complete all of this, however this extra work can save some time when there is an emergency or conflict.

5.2 BCP 38/84 for BGP

BCP 38 [18] and BCP 84 [6] are Best Current Practices to apply ingress and egress filtering (i.e., filter for ingoing or outgoing traffic). These two BCPs describe how filters can be applied, especially focussed on preventing DDoS attacks where spoofed addresses are used. These techniques can also be used on edge routers, and/or core routers depending on the network setup, to discard any unauthorised route announcements where the prefix can not originate from that certain neighbour.

If one wants to apply these techniques in his network, it is really important to get to know the peers. For filtering it is really important to know what kind of routes to expect from the peer and to know how to reach the neighbour when his filters needs to be updated. All of this has already been discussed in 5.1. For now, the most important thing is to know the prefixes that one wants to announce and the prefixes that one expects from its neighbours. If all of this is clear, one can configure its ingress and egress filters to only let some prefixes through and to only announce the prefixes that have been delegated to him. When one configures the ingress and egress filtering, also pay attention to the length of the prefixes that one wants to accept. Problems that might occur if this is not correctly done can be seen in Section 5.4.

It is also important that one implements the techniques described in BCP 38 and BCP 84, in order to prevent source spoofing. During a mail conversation with Rob Vercouteren, he stated that with today's hardware the overhead with route filtering is minimal. This used to be a problem, but due to more powerful hardware this is not a problem any longer. The main reason why route filtering is not implemented by all network operators is the required work involved in having the

filters up-to-date. Therefore, it would be easier if there is another method to be assured that only valid route announcements are accepted. The following sections will discuss some of these, but unfortunately they are not mature enough to solely depend on these solutions. It is therefore recommended to have appropriate filters in place.

Configuration example route filtering

In Appendix B configuration examples can be found for ingress and egress filtering of route announcements.

5.3 Maximum routes accepted

If it is impossible to implement ingress and egress filtering of route announcements in one's network, it is advised to have at least implemented the basic security measures. This section, Section 5.4, and Section 5.7 will outline some of these basic security methods.

In this section it is outlined how the security of one's network can be improved by configuring the maximum routes that one accepts from its neighbour. Some network operators may know it as the *maximum-prefix* option on Cisco devices. This option provides some basic security features, since one only specify how many routes one is willing to accept, which from our point of view is just damage control. This security measure defeats an attack where an attacker announces a lot of routes (e.g. an attack where an attacker wants to hijack a lot of routes). However, when one wants to implement it, it is really important to know beforehand how many routes are advertised by the peer. It is therefore important as well to maintain good contacts, by means of email or other preferred methods, when the peer wants to announce more routes when more prefixes have been allocated to him.

This is trivial to implement and therefore is not a burden for network operators to maintain. However, it could become a problem when a peer frequently changes the amount of routes it wants to announce. It is therefore important to have good contacts with the peer and talk about problems that might arise if he or she causes a problem with this limitation.

Configuration example maximum prefixes

In Appendix C configuration examples can be found for configuring the maximum amount of prefixes to accept.

5.4 Max prefix length

Depending on the implementation of one's router, it is important to specify the maximum prefix length[1]. This can prevent an attack where an attacker injects a route that is more specific than the originally announced route. Once again, it is really important to have good communication channels with the neighbours, since it can also have a legitimate reason to announce a more specific route[2]. In this case, communication channels are referred to as channels to update one's neighbour of the routes one wants to announce.

[1] This may be known as the *prefix-length-range* on Juniper devices

[2] It can be used to redirect the traffic to a scrubbing centre.

It could however become a problem when someone wants to legitimately, for mitigating a DDoS attack, announce a more specific than is allowed by the configuration. The route would not be accepted, however injecting more specifics is a frequent method used by attackers to hijack prefixes as outlined in [58] and [59]. Therefore, it is important that there are adequate prevention mechanisms in place. However, when someone wants to announce a more specific for a legitimate reason, it might be the case that this route is not accepted. Most network operators, as outlined in [58], will not accept any prefixes that are more specific than a /24. However, this is only the beginning, tight filtering would be more adequate in the case where one is always announcing a /22. An attacker can still announce a more specific, in this case a /24, and it will be still accepted if network operators only check if the prefix is not more specific than a /24. Therefore, it is recommended that there are filters in place that really check the route, instead of one general filter that checks all routes. However, administering these filters can become a burden when peers change the prefixes they want to announce and how specific they are, which could be a reason for a network operator not to implement this. Once again, we want to emphasise the importancy of having adequate filters in place, as it could otherwise happen that one is accepting a route that might be hijacked. If this is the case, the traffic could be eavesdropped or terminated at the attacker's side.

Configuration example maximum prefix length

The configuration examples are the same as for Section 5.2. Therefore, the reader can take a look at Listing B.1 and Listing B.2 in Appendix B.

5.5 IRR

One other method of creating filters is using the Internet Routing Registry databases (IRR) databases. These databases hold information about routing on the Internet. There are tools [43], [43], [12] to automatically query the IRR database and create filters based upon them. The tools mentioned before can provide one with the configuration for various platforms. Therefore, it is crucial that the information in the IRR databases are up-to-date and reflect the routing policies of the peers that one wants to apply filtering on. Various research has been done to check if the IRR databases are current and the results of the various reports are somewhat different [52], [60], and [25]. However, if filters are based on IRR and some company has malicious intentions, it is possible that it adds a route object to the database. Therefore, an additional method is required to check if the AS is authorised to announce the network. This method is known as BGP Origin Validation and will be discussed in Section 5.8. We suggest to use these tools that create filters based on IRR information if, and only if, one is assured that the IRR routing information of the peers is consistent and reflect the peer's current routing policy. If this is the case, it could be a great tool to automatically create filters based upon this data. One additional recommendation is to automate this process, the process of generating the configuration files and pushing it to the routers, since it will happen that peers update their routing policies. If it is not feasible to automate this process, one needs to be assured that the peer can notify the network operator of one's network when the routing policy changes (e.g. in the case when the peer announces new routes). To be assured that the peer knows how to notify the network operator, it is important to keep the contact information up-to-date as discussed in Section 5.1.

5.6 AS_PATH filtering

One other additional security measure, is filtering by AS_PATHs. This can prevent scenarios where a route is still being announced by the same neighbour, only the route now originates from somewhere else. When one has implemented ingress and egress filtering, this attack would still be possible, if and only if there is an ingress rule for this prefix of this specific neighbour. When AS_PATH filtering has been correctly implemented this attack would not be possible, because the AS_PATH filter looks at the AS numbers in the AS_PATH attribute. If this attack is executed, it would generate an extra AS number in the AS_PATH attribute. With correct AS_PATH filtering in place, the announced route will not be accepted, because there is an additional AS number in the AS_PATH.

AS_PATH filtering can be applied, in combination with other techniques, to prevent route leaks and/or routes that have an unexpected AS_PATH. However, creating various filters can become quite hard to manage if one has to do it for all its neighbours. However, various clever regular expressions can be used to check for certain ASes, which can minimise the number of filters one has to implement. As with any other filter, small mistakes can harm one's network. Only in this case, where an AS is mistyped in the filter, it could occur that all routes from that AS would be accepted or rejected. Therefore, special attention has to be paid when configuring these filters.

Configuration example AS_PATH filtering

In Appendix D configuration examples can be found for configuring AS_PATH filters.

5.7 Securing the BGP session

During the interview with KPN in Section 4.1.1, it was outlined that trust is the biggest issue on the Internet. This raises the question if one can really trust its peer and if one is assured that he is indeed talking to the intended peer. It could be possible that traffic, especially when one makes use of BGP multihop, is modified along the way or is being tapped. The biggest problem that needs to be solved is that one needs to be assured that he is talking to the intended peer and not to somebody else. One of the solutions is to use MD5 for authentication. However, there is also a possibility to use IPsec, which solves the shortcomings of MD5 in BGP as explained below, to secure the BGP traffic.

> BGPv4 md5 [...] authentication provides mutual authentication between the local and remote peer at connection origination, and secures session and data traffic on a per packet basis but, does not provide per-packet authentication, integrity, confidentiality or replay protection, therefore leaving the BGPv4 peering session vulnerable to attack. In addition, BGPv4 peering authentication, [...], does not provide for a protected cipher-suite negotiation. Therefore, BGP md5 authentication provides a weak security solution. (From: [62])

IPsec tries to address the above shortcomings and protects the BGP session and data. However, there are some problems with the use of IPsec for securing the BGP traffic. The main problem is the configuration that is required to make it work. One additional problem are possible re-keying issues, like the one with MD5 authentication, where one needs to exchange the new keys with

the neighbours. This allows room for errors and it could be possible that the BGP sessions are terminated, which may happen on an unforeseen time, this may result in unexpected downtime.

However, additional methods exist to mitigate BGP session hijacking attacks or TCP reset attacks. If possible, one can use private addresses for the peerings or use addresses that are not reachable anywhere else. This would prevent the TCP reset and BGP session hijacks, since they can not reach the router from the Internet. This solution however, is not possible when one wants to use BGP multihop, except when there is some tunnel in place and do the BGP session in that tunnel.

Configuration example MD5 and IPsec

In Appendix E configuration examples can be found for configuring MD5 and IPsec.

5.8 BGP Origin Validation

Resource Public Key Infrastructure (RPKI) [29] is a PKI system that supports improved security of Internet routing. This framework is based on X.509 PKI certificates with two extensions. One for binding a list of prefixes to the subject of the certificate and the other of binding an AS number to the subject of the certificate. The structure in which the certificates are made, mirrors the way prefix allocation is done by the Regional Internet Registries (RIRs). Route Origination Authorizations (ROA) states the AS number that is authorised to originate a certain IP prefix. This ROA is then signed by the private key of the resource holder, which creates a chain of trust. This offers validatable proof of holdership of the prefixes. Origin Validation is based on these ROAs and are used to check if the originator of the route announcement is authorised to announce the prefix.

However, this system can be circumvented. It is possible to do a path shortening attack, and is described in [20] as shown below.

> The second is a path-shortening attack in which an attacker announces a short bogus path to a prefix that terminates at the authorized origin AS. (From [20])

To make it more concrete, the AS_PATH attribute, which is used to get the origin AS, is not protected. It is therefore possible that someone with malicious intentions announces a route, with the origin AS as the AS that has the holdership of the prefix. The validation will succeed and without ingress and egress filtering, as described in Section 5.2, the route announcement will be accepted and may possibly be announced to other peers as well[3]. Without something in place that protects the AS_PATH attribute from being modified, this attack would be possible. Fortunately, work has already been done to protect the AS_PATH in the project named BGPsec. BGPsec will be discussed in Section 5.9. This could be one of the reasons to not implement Origin Validation, since it can be circumvented. There is also a problem with the ARIN RPKI Relying Party Agreement that causes network operators not to implement Origin Validation, because it is required that the agreement is signed before using the ARIN's Trust Anchor Locator [3]. It is necessary to have access to this Trust Anchor in order to validate route's in its region. There is also the question if network

[3]This depends on the configuration of the peer accepting this route announcement and the peers it is connected to.

operators can trust the trust anchor and if it is properly secured. If the trust anchor is compromised and all the ROAs are mangled with, such that validation fails, it can have a devastating effect on the Internet, since routes will not be accepted when verification fails.

However, we still recommend using Origin Validation to validate the routes. With BGPsec coming, and when it is used in combination with BGPsec, one has a lot of protection in place that makes it hard to accept rogue routes.

Configuration example BGP Origin Validation

In Appendix F configuration examples can be found for configuring BGP origin validation.

5.9 BGPsec

BGPsec provides protection for the AS_PATH attribute, which is not provided by BGP Origin Validation. Unfortunately, it is still in development and the only implementation found so far is in BIRD [56]. A good explanation on how BGPsec cryptographically assures that the AS_PATH has not been tampered with when validation succeeds, is shown below.

> The BGPsec framework proposed for securing the AS Path also makes use of a local RPKI cache, but it includes an additional element of certification. The additional element of the security credentials used here is an extension to the certification of AS numbers with a set of operational keys and their associated certificates used for signing update messages on eBGP routers in the AS. These "router certificates" can sign BGP update attributes in the routing infrastructure, and the signature can be interpreted as being a signature made "in the name of" an AS number.
>
> In the BGPsec framework, eBGP speaking routers within the AS have the ability to "sign" a BGP update before sending it. In this case, the added signature "covers" the signature of the received BGP update, the local AS number, the AS number to which the update is being sent, as well as a hash of the public key part of the router's key pair used to sign route updates. The couplet of the public key hash and the signature itself is added to the BGP protocol update as a BGPsec update attribute. As the update traverses a sequence of transit ASes each eBGP speaker at the egress of each AS adds its own public key hash and digital signature to the BGPsec attribute sequence.
>
> (From [21])

When a route has been propagated through the network, and every intermediate router signed the update, a whole chain of digital signatures exists. Due to this whole chain, it is possible to validate that the corresponding signature of that node was correctly generated on behalf of that AS in the AS_PATH. As stated above the signature covers, among other information, the local AS number and the AS of the peer that the update is destined for. This prevents the man-in-the-middle attack, where a legitimate announcement destined for a specific peer is send to another peer. Therefore, if every signature can be validated of the announced route, it can be cryptographically assured that no-one has tampered with the AS_PATH.

Since BGPsec is still a draft, it will take a while before it is implemented. If BGPsec is finalised and network operators start to implement it, it could drastically improve the security of BGP and therefore the routing on the Internet. However, as with most protocols, it will take time before

network operators start to adopt this new protocol. To get the most out of BGPsec, such that routes with a mangled AS_PATH would not be accepted by any network, everyone should implement it. If network operators do not implement BGPsec, in combination with Origin Validation, it would still be possible that rogue routes will be accepted by their router.

We advise network operators to stay updated about the latest developments of BGPsec and start experimenting with BGPsec in an early stage. This will allow network operators to integrate BGPsec more quickly in their network, since they gained some experience with BGPsec during the trial period.

Chapter 6

DDoS Analysis

Denial-of-service (DoS) attacks or distributed denial-of-service (DDoS) attacks are attempts to make computer machines unavailable to the Internet users. DDoS attacks can be performed on a network level, and on application level. Network-level DDoS attacks utilise network bandwidth that might otherwise be used for legitimate traffic. A lot of DDoS attacks are based on the fact that address spoofing is possible on computer networks. One of the biggest and most well-known attacks where source spoofing is used, are amplification attacks. Another example of a DDoS attack is the SYN flood attack, where a lot of TCP connections are initiated by a SYN, using a spoofed source address. This results in the server responding with a SYN-ACK to the falsified source address. If that actual source address will not respond with the ACK message and will not terminate the connection, it will leave the TCP three-way handshake unfinished. For every TCP handshake the system will allocate memory. When there are a lot of unfinished handshakes, a system will eventually be unresponsive to any other request. The application-specific DDoS attacks are becoming more of a threat, because these attacks abuse the behaviour of protocols like TCP or HTTP. Also, the impact of DDoS attacks are becoming more problematic, because a lot of people are dependent on Internet services (e.g. banking, governmental services, etc.). A solution for mitigating the DDoS attack would be to send the traffic through a scrubbing server. However, it is possible that the attack has a larger volume than the scrubbing server can handle. From the company's point-of-view, they could lose a lot of money (e.g. when banks can not fulfil their money transaction). Or hackers that initiate DDoS attacks and will stop when business pay a large amount of money. Finally, sniper attacks are becoming more of a threat as well. These attacks happen during a DDoS attack, where the attacker hides the attack in the high volume of DDoS traffic. This kind of attacks are hard to differentiate from the DDoS traffic, as they might have similar characteristics as the DDoS traffic.

According to a recently released paper by Kaspersky [24], the cost of DDoS attacks for small-to-medium-sized businesses per incident are on average $52,000. This costs include all lost business and reactive IT spendings. For larger enterprise businesses, the costs for DDoS attacks per incident are on average $444,000.

6.1 Scrubbing

A technique that is commonly used for DDoS mitigation is the use of scrubbing. When a company detects a DDoS attack, the malicious traffic will be redirected to special mitigation server to filter out the malicious traffic, leaving the company unaffected. These scrubbing machines apply DDoS filtering and routing techniques to reduce the DDoS traffic to acceptable limits [49]. Scrubbing devices use sampled data as input. Using this input, the scrubbing tool will propose a countermeasure based on attack signatures that the operator can use. A few of the companies offering those machines are Arbor Networks, Cisco, and Juniper. Besides dedicated scrubbing machines, there are also cloud solutions for scrubbing in the form of scrubbing centres. Companies like CloudFlare and VeriSign offer these cloud solution, where a company does not have to own the scrubbing equipment themselves. Especially for small businesses that do not want or can buy their own scrubbing

machines, cloud scrubbing can be very useful. In The Netherlands, a few companies bundles their powers and created the Nationale anti-DDoS Wasstraat[1] [39].

As mentioned before, scrubbing is one of the first mitigation techniques used to mitigate a DDoS attack, and is used by a lot of companies all over the world. For most DDoS attacks scrubbing would be a good solution. However, it is not always the right solution. When cloud scrubbing providers does not have the capacity to mitigate DDoS attacks for more than two companies at the same time, one company will stay in the dark.

6.2 BCP 38/84 for DDoS

Up until now, we only covered the filtering of route prefixes that one sends and receives using BGP. As described in BCP 38 and BCP 84 it is important to drop packets that have a spoofed source address. This requires some sort of filters or Access Control Lists (ACLs).

The general idea behind the ACL is to inspect every packet and verify if it is expected that such a packet arrives on that specific interface, with such a source address. If the filters have been properly configured, then packets with spoofed source addresses will be dropped. However, in large networks it could become a cumbersome process, since this needs to be done manually.

In case of an ISP where they have to add, remove, or change routes a lot, configuring ingress and/or egress filtering requires a lot of manual labour and is the reason why it is not always used. It could prevent a lot of DDoS attacks, since most of the time attackers use spoofed addresses. This requires that the access lists stay up-to-date, in order to prevent legitimate traffic from being dropped. There is a solution for this problem that does not require this manual configuration, which is discussed in Section 6.5.

Configuration example source filtering

In Appendix J configuration examples can be found for filtering based upon the source address of a packet.

6.3 Intrusion Detection Systems

Intrusion Detection Systems (IDSs) are able to scan the incoming Internet packets for threats on the network. There are signature-based IDSs, that use a signature database to compare the network packets for known malicious threats. This is similar to how anti-virus software works. Besides the signature-based IDSs, there are also anomaly-based IDSs. This kind of Intrusion Detection Systems compare network traffic against the normal network use (use of bandwidth, which protocols are used and which devices are connected to each other) [48]. A signature-based IDSs might not be able to detect today's DDoS attacks. Today, IDSs need to be a combination of both anomaly-based, and signature-based systems that are able to detect today's sophisticated DDoS attacks [8] [53].

IDSs will generate a lot of false positives right after the implementation, and need some extensive tweaking, as mentioned in Section 4.4. Unfortunately, Intrustion Detection Systems can also be victim of DDoS attacks itself, when these systems are placed like a firewall in a network, so traffic passes through them. In this way the IDS will get overloaded in the case of an incoming DDoS

[1] National anti-DDoS Scrubbing Centre

attack targeted at someone behind the IDS, since all the network traffic flows through the IDS [48]. An IDS is therefore not a mitigation solution, but could be a helpful tool that provides specific attack flow information. It requires a specialised team of analysts that are able to use the tool and are able to interpret the data [46]. As mentioned before in Section 4.4, a Security Information and Event Management (SIEM) system is very helpful to correlate event data from e.g. system log files, and Intrusion Detection Systems. This kind of tool helps to find for instance DDoS attacks with sources like router and/or server application logs, and the alerts generated by IDSs.

Intrusion Detection Systems need to do Deep Packet Inspection in order to look into the network packets and find malicious threats. When packets are encrypted, Intrusion Detection Systems will not be able to compare the packet content to the signature database, or check for anomalies against the baseline of the network. A major risk of signature-based IDSs is that one need to keep the signature database up-to-date. When a newer signature is available, and not installed, an signature-based IDS will not be able to detect that particular threat [53].

Be aware that an IDS can not be placed everywhere and that a user's privacy should be respected. There are laws that might prohibit one from looking into the traffic.

6.4 NetFlow

A way to be able to analyse the network flow after an attack has taken place, is to use NetFlow data. NetFlow captures network flow data from routers in a network and sends it to a NetFlow collector. Network administrators use the data to find anomalies in the network flows and has been proved to be valuable [50]. For instance, network administrators use NetFlow to find the entry point of an attack in a network by backtracking the attack from each router in the network. Besides analysis after an attack has happened, NetFlow data can also be used to scan real-time for DDoS signatures by a DDoS analyser. Based on this information, DDoS traffic can be send to a DDoS scrubber to mitigate the attack. Besides NetFlow, there is also a newer standard named IPFIX [9]. This standard is made by the IETF and is based on Cisco's NetFlow v9. The IPFIX specification has some extra features, like defining templates for dynamic data definition [36]. Besides NetFlow and IPFIX there are also NetFlow equivalents (e.g. sFlow, Jflow, and NetStream).

The reason for implementing NetFlow within a network is that network administrators can investigate the network flows after they found an incident, which they want to investigate. Without NetFlow, or a NetFlow-like solution, it would be impossible to investigate the network flows. Only packet capture could replace NetFlow, but this would require a lot of disk storage, and packet capture is not always allowed by law.

One large risk of using NetFlow is the risk of being buried with too much data. When having too much data, one can choose to sample the data, e.g. one sample every five minutes. However, when the sample rate it is too high, it is possible that important flow data will not be included in the flow storage.

6.5 Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (uRPF) is a technique to block Internet packets that use forged source addresses. For uRPF to be effective, it would have to be implemented in front of every potential attack source [46]. Every packet that arrives on an interface and that should not use that specific interface for that source address, are considered forged and will be dropped. In uRPF there

is a loose mode and a strict mode. In loose mode, the only check is whether the Internet packet has a source address with a corresponding prefix in the routing table. The router will not check whether the interface expects to receive a packet with that specific source address of the received packet. If a corresponding prefix is not found, uRPF loose mode will not accept the packet. In strict mode it will check if the interface expects to receive a packet with a particular source address prefix [22]. In theory, when the router maintains a Forwarding Information Base (FIB), uRPF in loose mode might not help. In that case all traffic will be accepted, because the router has all the routes within the routing table. This might be one of the reasons why ISPs do not implement this [1].

uRPF example

In Appendix G configuration examples can be found for configuring uRPF.

6.6 BGP FlowSpec

BGP FlowSpec is a relatively new protocol/tool that can be used to dynamically mitigate DDoS attacks. BGP FlowSpec is an extension to the traditional BGP protocol that enables the communication of filters through the existing BGP network. This allows network operators to place filters in their own network and in their neighbours' network. For example, what traffic needs to be redirected to a scrubbing system, what packets needs to be dropped, or what traffic needs to be rate-limited [31]. The reason for using BGP to propagate filtering information is that DDoS mitigation actions are getting distributed in one's network and the neighbours' network, and filtering will be performed closer to the source of the attack. Since BGP is probably already used in one's network and thus already have 'trust' in one's peer, it makes sense to distribute the FlowSpec routes using BGP.

As mentioned before, FlowSpec is very new, and not all routers will support it, especially the older models. Therefore, a lot of companies are not able to implement it, and will wait until it has been tested with higher volumes of network traffic (more than 300 Gigabit per second). Since it is new, it would be likely that people make configuration errors. For instance, it is possible that when validation is not correctly configured, a peer advertising a filter to one's network will block traffic for another peer. It is an interesting technique that looks promising and we recommend network administrators to stay informed about the developments.

6.7 Trusted Networks Initiative

The Trusted Network Initiative is a non-commercial initiative to create a last-resort measure against DDoS attacks. As mentioned before in Section 4.3, TNI is not a permanent solution for DDoS attacks, it does not solve the problem completely. TNI is created with the idea that scrubbing does not always have the desired result. Contracts with scrubbing companies are relatively expensive and when the limit of the contract exceeds, one would have to pay extra. Sometimes the scrubbing companies do not always have the capacity to scrub the network traffic of multiple companies at the same time, so they have to prioritise on which company to scrub first, leaving another company in the dark. DDoS attacks are getting more advanced, taking much longer than before, and the

attacks are implemented in a smarter way. The need for temporary or permanent solutions is very high.

Where other concepts try to invent something completely different, TNI wanted to use techniques that are available today, with the use of strict TNI policies [61]. In this policy every peer of the Trusted Network Initiative make some promises about the technical measures they have to implement, e.g. the implementation of anti-spoofing measures like BCP 38/84, and that these companies have to operate a capable CERT and/or CSIRT [61].

The technical implementation of TNI is relatively simple. Every company that participates in the Trusted Network also maintains the original peering implementation. Additionally, these companies will also have an additional peering session on VLAN 112 to the route server of the Trusted Network, with a lower local preference than the normal peering. The Trusted Network is located at the transit networks AMS-IX and NL-IX on the symbolically VLAN '112', a wink to the Dutch and European emergency service number. When a connected company suffers from a DDoS attack and is not able to mitigate the attack using the traditional methods, the company can decide to temporarily disconnect the route to the normal Internet and completely rely on the trusted network. In this way the participants of the trusted network are still able to visit the network of that company. When large ISPs in The Netherlands participate and for instance a large Dutch bank, customers will still be able to visit the website of that bank via the Trusted Network peering. However, one will lose a part of the Internet users, because only users that are connected via the trusted network are able to access that website or service.

The Trusted Network Initiative does not try to create a small safe island within the Internet. Companies can decide for themselves if they want to 'raise the Internet bridge' of the public Internet, thus disconnecting from the normal Internet. Therefore, this concept does possibly not conflict with the Net Neutrality laws within The Netherlands. It is the choice between not being available on the Internet at all, or temporarily for a small amount of people. It is a grey area in the Net Neutrality Law.

TNI can be used on a temporary basis, but a company can also decide to use the trusted network permanently. This is still a debate within the Trusted Network Initiative, because especially ISPs are not that happy about this. Normally, ISPs do not like to peer directly with other parties, because BGP peering is based on equality. It would also be possible for ISPs to do private paid peering. However, we would recommend connecting to the TNI route server to be able to exchange traffic between every other peer in emergency situations.

In the Czech Republic there is a similar project, called The Fenix Project [55]. The main difference between the Dutch Trusted Networks Initiative and the Czech Fenix project is that in The Netherlands there are two Internet exchanges that participate versus one Internet exchange in the Czech Republic. When TNI is out of the trial phase and when the project has attracted more participants, it is possible that TNI will start connecting to international project to create a large(r) trusted network.

The implementation of TNI is tested with a Proof of Concept with the parties NLnet, NL-IX, and SURFnet. SURFnet created a DDoS attack and NLnet disconnected the normal BGP connection to the Internet, and was still available through the trusted network. At the moment, the Trusted Networks Initiative is in the trial phase where they are fine-tuning the concept, and are going to establish a Trusted Networks Initiative Foundation. The next phase would be to launch the concept in its final form, and to attract more participants, nationally and possible internationally.

The Trusted Networks Initiative does not solve the source of the DDoS problem itself. In our

opinion the Trusted Networks Initiative is an interim solution that gives us some breathing room in order to find a permanent solution. A permanent solution would be that everyone implements the proper security measures. However, this is opportunistic idea that takes time, because everybody connected to the Internet should implement the before mentioned security mechanisms. We do not see this happen within the upcoming years. That is why we need an interim solution, in order to run our critical infrastructure during an emergency. The Trusted Networks Initiative provides us with such an intermediate solution, and we would recommend network operators that serve critical infrastructure to participate in this initiative, and only make use of it when there is an emergency.

Chapter 7

Phishing Analysis

Spam emails are not really an issue any longer, because spam filters are working fairly well these days according to the abuse specialists we have spoken to, as mentioned in Section 4.1.3 and Section 4.4. Phishing on the other hand is one of the largest problems with email these days.

> Phishing is a kind of attack to steal sensitive information, e.g. usernames, passwords, banking details, by masquerading as a trustworthy entity in an electronic communication [e.g. emails, websites, etc.]. (From: [51])

These days, it is very easy for spammers and phishers to abuse the email system. Back in late seventies and early eighties when email was invented, there was almost no security and integrity of emails implemented, since the amount of email users was very low. When the amount of email users started to grow, the amount of abuse started to grow as well. Up until now, there is almost no verification of the integrity of emails.

For the average user it is very hard to distinguish legitimate websites or emails from the phishing websites or emails. In research from 2006 [11], the best phishing website was able to fool more than ninety percent of the participants. In The Netherlands there are a lot of fake emails from "banks". The phishers try to make users to click a link and try to make them fill in their username, password, and/or banking details. These emails seem to be legitimate, because these emails are mailed from, e.g. 'info@[name of bank].nl', or 'noreply@[name of bank].nl'. However, these mails are not really sent from a bank. These mails do most likely have a forged "MAIL FROM" address in the envelope and a forged "From" address in the header.

Listing 7.1: Example Mail header

```
Message-ID: <[removed from document]@os3.nl>
Date: Fri, 23 Jan 2015 11:20:46 +0100
From: Koen Veelenturf <koen.veelenturf@os3.nl>
User-Agent: [removed from document]
MIME-Version: 1.0
To: Wouter Miltenburg <wouter.miltenburg@os3.nl>
Subject: Test Mail
Content-Type: text/plain; charset=utf-8
Content-Transfer-Encoding: 7bit
```

7.1 Sender Policy Framework

With the use of the Sender Policy Framework (SPF) [26], and DomainKeys Identified Mail (DKIM) Signatures [10] it is possible to verify the origin of the email and check the integrity of the email. SPF uses a special DNS TXT record in which is specified from which servers emails are allowed to be send from for that specific domain name. The domain name in the "MAIL FROM" from the envelope, or "HELO" during the SMTP communication, is used to identify the domain name used by the email sender.

Listing 7.2: DNS SPF example (Google Mail)

```
kaveelenturf$ dig _spf.google.com TXT
[...]
_spf.google.com.          300       IN        TXT       "v=spf1
   include:_netblocks.google.com include:_netblocks2.google.com
   include:_netblocks3.google.com ~all"
[...]
kaveelenturf$ dig _netblocks.google.com TXT
[...]
_netblocks.google.com.    3600      IN        TXT       "v=spf1
   ip4:64.18.0.0/20 ip4:64.233.160.0/19 ip4:66.102.0.0/20
   ip4:66.249.80.0/20 ip4:72.14.192.0/18 ip4:74.125.0.0/16
   ip4:173.194.0.0/16 ip4:207.126.144.0/20 ip4:209.85.128.0/17
   ip4:216.58.208.0/20 ip4:216.239.32.0/19 ~all"
[...]
```

Using the example of the Google Mail's (SPF) TXT record in Listings 7.2, all the addresses within the 'netblock' ranges are accepted as source of emails send from the @gmail.com domain. The '~all' flag in the TXT record triggers the SoftFail mechanism.

> SoftFail: The SPF record has designated the host as NOT being allowed to send but is in transition, the intended action is to accept, but also mark. (From: [42])

As mentioned above, with the SoftFail mechanism, email that fails the verification will be generally accepted. The email will be tagged and it depends on the receiving party if the mail ends up in the spam folder or just in the inbox of the user. The SoftFail mechanisms is generally used for testing purposes. If the test results of the SPF filtering are acceptable, then one could choose the Fail mechanism, where email will be rejected if the source address of the email does not comply with the SPF record.

The reason why SPF is not implemented by some of the ISPs and mail providers is the amount of email servers a company can have. Over the years, KPN took over multiple ISPs in The Netherlands, namely Telfort, XS4ALL, and Planet Internet. All those ISPs maintained their own mail servers, and KPN took them over and carried on maintaining them. This results in a more complicated infrastructure for KPN in which they have to implement the email authentication mechanisms, as mentioned in Section 4.1.3. The initial implementation of SPF would take a lot of time, because the network operator needs to know exactly where every mail server is, and which IP addresses are allowed to handle the email traffic.

However, it is only the initial configuration that takes a lot of time. Whenever one installs a new mail server, that would be the only time one needs to edit the SPF record to add the corresponding IP address. This procedure has a low maintenance cost and we suggest that every ISP and mail operator should implement SPF, because it has a high impact on decreasing the amount of mail where spoofed email addresses are used.

Configuration example SPF

In Appendix H configuration examples can be found for configuring SPF.

7.2 DomainKeys Identified Mail Signatures

> DomainKeys Identified Mail (DKIM) defines a domain-level authentication framework for email using public-key cryptography and key server technology to permit verification of the source and contents of messages by either Mail Transfer Agents (MTAs) or Mail User Agents (MUAs). (From: [10])

The DKIM specification defines a mechanism by which emails can be signed, using a domain specific private key. Email recipients can verify the signature by querying the signer's public key by using the special DNS DKIM TXT record. The approach of the DKIM specification differs from message signing with S/MIME [19], and OpenPGP [7]. The signature used by DKIM does not appear in the body of the message, but is placed in the header of the message, as shown in Listings 7.3, so the average user will not notice the use of DKIM. In implementations like S/MIME one need a central trusted authority (CA). With DKIM there is no dependency on the public and private key being issued by a well-known authority, because the public key is distributed by a TXT record in DNS.

Using this public key, the email client can confirm that the email was sent by the owner of the private key for the signing domain. As shown in Listing 7.3, a signature is created based upon the private key of the sending email server. The *bh* tag is used for the hash of the canonicalised body. The *b* tag is used for the signature data. Both the *bh* and the *b* tags [10] are encoded in base64. Using the public key published in DNS, the receiving email server can confirm the integrity of the email content.

Listing 7.3: DKIM example: email source

```
Return-Path: [removed from document]
[...]
        for <koen.veelenturf@os3.nl>; Sat, 24 Jan 2015 16:19:55 +0000
            (UTC)
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple; d=[removed from
    document];
        s=mail; t=1422116395;
        bh=130coiLeCaS2/lMxA0nTPodUpSuyMYiA9lZW70NeduM=;
        h=Date:From:To:Subject;
        b=moXsUeMzZNJLCHJ1Aym8XxCYrURMnjLiUO1mqj9MDgrdK1iX4U8zrIqRO
            37jKOXhHslEu85fsAfAh9YJnXa8XaY37d9Dkug7GiohFvvxu9yY43rZnBV
            Z7gtlGTnMxf1eFrSGTIngt4dUu5tS5QbzcfV8d18y3tL9CtqPZXSy/c=
Message-ID: [removed from document]
Date: Sat, 24 Jan 2015 17:19:55 +0100
From: Koen Veelenturf <[removed from document]>
User-Agent: [removed from document]
MIME-Version: 1.0
To: koen.veelenturf@os3.nl
Subject: Test with DKIM
Content-Type: text/plain; charset=utf-8
Content-Transfer-Encoding: 7bit
```

```
X-OriginalArrivalTime: 24 Jan 2015 16:19:57.0142 (UTC)
    FILETIME=[980D9F60:01D037F1]
X-RcptDomain: os3.nl
```

When network operators implement DKIM on their mail servers and they turn on the checking for DKIM, phishers will not be able to send email 'on behalf of' email address x. Phishers will not be able to edit the content of an email with, for instance a man-in-the-middle attack, because they do not have the private key for signing the email.

The reason why DKIM is not implemented by the majority of ISPs and mail providers is, just like with SPF, the amount of email servers a company can have. Besides the amount of servers, DKIM will also take an additional load on the mail server, because every email needs to be signed with the server's private key. Every domain needs to have its own private key, and needs to publish the corresponding public key in the proper DNS records. If an operator maintains multiple domains, setting up DKIM could be a time consuming process.

It is very important to implement SPF, because SPF decreases the amount of mails with spoofed source addresses. DKIM is a nice addition to SPF, in order to verify the integrity of received emails. Therefore, SPF would have a higher priority for implementing it on mail servers, so that at least mail servers that do verify SPF records do not let phishing mail with a spoofed address through.

Configuration example DKIM

In Appendix I configuration examples can be found for configuring DKIM.

7.3 Domain-based Message Authentication, Reporting & Conformance

The SPF and DKIM frameworks offer the possibility of filtering out email that is not send from the authorised mail server. Domain-based Message Authentication, Reporting & Conformance (DMARC) offers a standard on how email receivers perform email authentication using the SPF and DKIM mechanisms. The technical specification helps to reduce the potential problems of operating, deploying, and the reporting related to SPF and DKIM [13].

DMARC reports contain feedback, aggregated and forensic reports from mailbox providers on every email that did not pass the authentication. The reports are generated in XML format. At the moment there are no open-source solutions for parsing the XML files into a human-readable format. At the moment the commercial dmarcian [14] is one of the few tools for DMARC information.

> Nearly two billion email accounts worldwide are protected by DMARC. [...] Outlook.com reports a 50% drop in reported phishing in 2013, in part due to DMARC, and more than 25 million email messages spoofing PayPal were rejected during the 2013 holiday buying season. (From: [13])

In case of Google, Yahoo, and Microsoft's Outlook.com (formerly Hotmail) made the decision to implement it to reduce spam for their users. For these companies DMARC is working fairly well, as mentioned above. If mail services like Google Mail, and Outlook.com decide to force the checking of SPF records, companies that do not have a SPF record, will not be able to send email to those mail services. Since Google Mail, and Outlook.com have millions of users, companies will

almost be forced to implement SPF. At the moment, KPN is implementing DMARC on the first mail server and other servers will follow this year. The implementation by large ISPs would result in a large decrease of spam and phishing email.

We recommend companies to implement SPF according to the DMARC specification. If time and money allows, we would also suggest to implement DKIM, so that other's can verify the integrity of the sent emails. Finally, DMARC will help to monitor the use of the email addresses, albeit legitimate email or illegitimate email.

Chapter 8

"No business case"

All of the previous sections showed how one can make a network more resistant against certain types of attacks. However, to implement all of this one will need time to get familiar with the subjects and will cost money. Therefore, it is important to have a good business case that allows one's manager to understand the risks and potential losses for the company. Most of the time businesses do not want to lose face. Therefore, it is important to emphasise this in the business case. This chapter will provide information to help one to create such a business case.

8.1 Return on security investment

It costs money to improve the overall security of a network and it costs money as well to keep it safe. It might not sound really attractive for a manager to invest money in one's network, but it might actually safe money in the future. The best way to convince a manager is in terms of money. How much will the company lose when they are offline for a day? What will it cost to provide support to the customers who might call the help-desk? What are the other risks? What will it cost the company when personal identifiable information is copied and made available on the Internet? What will happen when proprietary technology is made available on the Internet? If a company has multiple SLAs for its customers, what will happen when the company can not meet these SLAs? There may be some additional costs that are hard to quantify. In the case of reputation damage it is not really possible to calculate the losses on beforehand, since one can not really expect how customers would react on a certain incident. For example, if private information of a customer has been compromised, it might be the case that the customer will break the agreement with one's company. All of this needs to be accounted for in the calculation of the company's losses. A great exercise to get familiar with estimating the costs of a security incident is [16]. It is advised to keep track of the incidents and the costs involved with the resolution of each of these incidents. This can help one to quantify the potential loss and to assist management decisions.

The other part is calculating the costs of implementing a certain security measure, including the additional support costs and the systems that might need to be replaced. If it turns out that it is much cheaper to let one invest time and money to secure a platform, rather than acknowledging the fact that there is an exploit in one of the systems and not doing anything about it, it could help to convince one's manager. In the interviews with KPN and the multinational company, which can be found in Section 4.1.1 and Section 4.4, it is shown how important it is to justify and to really convince one's manager of the business case.

8.2 Law

Depending on the laws that are applicable for an organisation, it could be the case that some extra security facilities need to be implemented in one's network or systems. For example, The Netherlands has the *Wet bescherming persoonsgegevens (Wbp)*[1] and *Telecommunicatiewet*[2], which

[1] Personal Data Protection Act
[2] Dutch Telecommunications Act

states that personal data should be handled carefully and extra security mechanisms should be implemented to protect the data. The *Telecommunicatiewet* is especially focussed on ISPs and state that data leaks should be reported to the *College bescherming persoonsgegevens (Cbp)*[3] of The Netherlands.

There are no laws that define how one should do routing and how to secure the routing. If one needs to convince a manager to implement BGP Origin Validation, it is better to convince him of the benefits of using BGP Origin Validation.

However, we highly recommend that one looks at the laws, if some security mechanism must be implemented by law or regulations, to secure the network. For example, ISPs need to take additional security measures, such that their customer data is protected.

8.3 Reputation

Reputation is truly important for a company[4] and a big security incident or even an emergency can lose a company face. If the company's reputation is damaged it could result in losses in terms of money and customers. It will take time before the reputation of the company is restored and the company should be aware of this. For example, think about the Sony hack [63] and imagine that this happens to one's company. Would the company lose customers and how long would it take to restore the reputation? This, and other cases are examples of what a hack can do to the reputation of a company. Make use of this and create awareness to the colleagues and managers about the consequences of a security incident or emergency.

8.4 Awareness

As was outlined in the previous section, it is really important to create awareness and that security is acknowledged as a business process. In the case of KPN, they improved their security a lot after they have been hacked as stated in Section 4.1.3. There is a lot of awareness in the company itself about the potential risks and that security is an important aspect of day-to-day operations. This is something one should aim for, to create awareness in the company about potential security risks. If more people acknowledge the fact that security is important, it becomes more natural to make security part of the business processes. Security should not be considered part of the 'operations' department in the company's hierarchy, but it should be very close to the CEO. In that case, the manager of the security team can directly talk to the CEO, if something needs to be decided or needs to have CEO approval. This can prevent cases that a project need to pass several managers, who have their own personal views on security.

[3] Personal Data Protection Commission

[4] Well not all as in the case of a bulletproof hosting provider or a transit that makes money out of DDoS attacks.

Chapter 9

Additional security measures

This chapter will outline how to add additional security measures into critical infrastructures. The items that are discussed vary from each other and begins with a more theoretical subject, and will end in a more technical part.

9.1 CERT

A Computer Emergency Response Team (CERT) is a team that consists of experts that handle computer security incidents. To be able to handle these computer security incidents it is important that information is shared between other CERTs. As has been widely discussed during the KPN interviews and the interview with the multinational company, see Section 4.1 and Section 4.4, the security of one's network depends on the information received from others, as stated by Rob Vercouteren and by Mark.

These CERTs exchange information on current attacks, send each other alerts about current security issues, and current activity about security activity. However, CERTs tend to help each other when there is a security incident, but this is again based on trust and one's reputation. In order to receive help, one needs to help others as well. Therefore, it is good to create relations between CERTs and to maintain them.

This report will not explain how to set up a CSIRT/CERT since this could be a dedicated report on its own. A good reference to start a CSIRT/CERT is [17].

9.2 Responsible disclosure

It is truly important to have a responsible disclosure policy and a company is really advised to implement this. If the reader's company did not implement a responsible disclosure policy yet, the reader is advised to talk to the manager to make time available for creating such a policy. By using a responsible disclosure policy one invites people to share weaknesses of the system with him. The policy does not provoke people or implies that people should hack one's system, it basically invites people to tell what is wrong when they find a weakness, without having a malicious intention with this information. As KPN explained during the interviews, it is a very effective method and it really helps KPN. Instead of not knowing where weaknesses in the system are, people tend to tell the company when they find them. Whereas previously people were afraid of having legal slumber, the company can now actually fix the weaknesses. When people can not tell a company that there is a weakness in the system, somebody else will find it, and this person might have a bad intention. A good guideline for creating a responsible disclosure policy is [40] and a good example is [54]. It is also important that a company needs to be able to process a responsible disclosure. This means a company needs to define its internal processes to deal with security findings. Companies risk losing contact with the discloser if the process is not handled properly or when the process is too complicated. The possible results could be that the discloser tells others about the weakness, like friends or a journalist, or changes its attitude and becomes a threat. It is important as well that

the policy is easy to find, otherwise people might not know how to report a responsible disclosure that correctly follows the policy.

Some companies have not implemented responsible disclosure policies, because they might be afraid that their company will be targeted by hackers. However, as KPN explained during the interviews, they only have benefited from it. We encourage companies to implement responsible disclosure policies, so that people know how to alert the company, and that there is a standard policy that can be used.

9.3 Monitoring

As told by KPN and A2B in Sections 4.1.2 and 4.2, it is really important to have monitoring in place. Without monitoring, it would be at least relatively hard to keep a network up-and-running and have a short response time in the case of an incident. Monitoring includes host software, but as well gathering traffic patterns of a networks using NetFlow, sFlow, or other protocols/software.

If possible, depending on the router, it would be useful as well to log and monitor the prefixes that the router receives. In case of network disruption, it could provide meaningful data for debugging purposes. It could happen that neighbours announce prefixes that they should not announce, and in the case of incorrectly configured filters, it could happen that traffic leaks to a neighbour. Something went already wrong here, since the filtering was not correctly done, but monitoring could at least minimise the duration and effect of the traffic that is leaked to a neighbour.

9.4 Patching

Patching should be done on a regular basis and as trivial as it may seem, it is relatively hard to patch all the systems in the network. It is relatively hard for a big company to find all the systems. This 'fragmentation' could pose problems when systems need to be patched, since it might be the case that a few systems were not patched by the patching system. It is advised that the infrastructure is built in such a way that continuity can be guaranteed when a system is patched with little to no impact. It is important to take special care with automated systems and be assured that all systems are included in this automation process. However, it could happen that not all systems are included in the automation process. Therefore, it is important to have a responsible disclosure policy, as outlined in Section 9.2. In case a system has not been patched and is vulnerable to a certain type of attack, people can responsibly disclose this vulnerability. If there is not a policy in place, people might tend not to report it, which can even pose a greater risk (e.g. someone with bad intentions that finds the vulnerability).

9.5 Asset management

As was pointed out during multiple interviews in Sections 4.1.1 and 4.1.3, fragmentation is a problem. At first sight this might sound like a trivial problem, but one can ask himself: "Do you know all your systems within your infrastructure?" A single weak spot in an infrastructure could have a devastating effect. Is there a solution for this in large ISP networks one might ask. Well, at least all of the suggestions from above, should be implemented. With monitoring it is possible to scan a network for outdated software and patch it. If possible, one should scan the configuration of

network equipment as well, look for known insecure configurations, and scan for unknown systems. Still, with all of this in place, there still can be an undetected system in the network. Therefore, one should have a responsible disclosure policy that in the case somebody notices this old system and sees a vulnerability, he or she can notify the operator.

Chapter 10

Conclusions

The Trusted Networks Initiative tries to create a last-resort measure against DDoS attacks when traditional mitigation measures do not help any longer. This is done by creating a network of trusted peers. Companies will maintain the current Internet peering, and they will also peer with the trusted network. When a trusted networks-connected company is under attack, this company can decide to 'raise the Internet bridge' of the normal Internet connection. Other companies that are connected to the trusted network and exchange routes with other peers, are still able to exchange traffic with the company that is disconnected from the Internet. At the moment, the Trusted Networks Initiative does not try to mitigate other common attacks on computer infrastructure. The principle of trusted networks does not solve the DDoS problem permanently. This measure will buy extra time to create more permanent solutions, like implementing anti-spoofing measures by all Internet companies.

The interviewed companies that maintain critical computer infrastructure are all suffering from the same range of attacks. However, not for every company the risks are the same. Therefore, it varies how the company prioritises the risk and how the network operator prioritises the risk. The interviewed companies identified DDoS as an issue, but phishing and BGP hijacking as well. However, it all depends on the type of company. Therefore, the risk of the identified problems for critical infrastructure may vary across companies, depending on what kind of company they are. One problem that most of the companies have is their business case for security and if the business simply 'cares' about their security. This is by far the biggest problem that companies suffer from, if there is no business case to implement a certain solution it will never be implemented. This results in leaving the company vulnerable.

In this report various security mechanisms are discussed that could prevent BGP hijacking, DDoS attacks, and phishing. For BGP hijacking it is basically filtering and Origin Validation, in combination with BGPsec, that would drastically help in securing BGP itself. However, Origin Validation should be used by everyone to make it really effective and BGPsec is still in a development phase. Therefore, network operators should filter and use Origin Validation if the business allows it. To prevent phishing and unwanted mail, one can use SPF and DKIM to filter mails and in order to prevent people from sending mails on one's behalf. However, it is the same with BGP Origin Validation, that everybody should do it before it becomes really effective. That is why we need to do filtering so that at least there is some protection mechanism in place in the meantime. For DDoS attacks, one could use the traditional scrubbing services, but that is only mitigating the problem and not solving it. To really solve it, all networks should do ingress and egress filtering such that it is not possible to use spoofed source addresses in a DDoS attack. This would drastically decrease the amount of DDoS attacks. The Trusted Network Initiative can help companies that need to exchange traffic with other partners, even in the case of a DDoS attack. It is not a permanent solution, but it allows network operators to solve the underlying problem and have a last-resort measure in the meantime. Other protection mechanisms have been outlined in the report as well, but it is up to the network administrator if this is possible in their network and if the business model allows it.

Having a solid business case is by far the biggest problem. A company might not see the

security risks that a network operator has identified. It is therefore important to know how to create a business case and convince a manager to invest time and money in applying the protection mechanisms. Various important protection mechanisms have been discussed and also why one should implement it.

The reason why techniques like anti-spoofing measures are not implemented is that the business has a higher priority than IT for a company. IT, and especially IT security, does only cost money and does not make any profit. It is hard to explain to managers why there should be invested in the IT security of a company. Therefore, it is very important as an engineer to create a good business proposal. In this proposal, one could argue the possible costs during an IT incident. Besides the money aspect, there are not a lot of laws and regulations on the topic of Internet security. However, in The Netherlands there is a law that forces companies to protect personal data, namely the *Wet bescherming persoonsgegevens (Wbp)*. This way, the Dutch government is regulating the security of personal data. However, there is also a law, namely the *Telecommunicatiewet*, that prohibits ISPs from filtering on their customer's network. Filtering could in certain cases actually help securing the network. It is a fine line between the privacy of customers and the security of infrastructures. Finally, there is not a lot of awareness with companies on topic of IT security. Therefore, it is very important to raise the awareness about IT security and to make IT security more important within the business process.

Chapter 11

Future work

The most important thing what is left to be done is to motivate companies to implement the suggested security mechanisms. As discussed in Chapter 5, the Internet will only be safer when other companies start to adopt the suggested security measures. Awareness and stimulation can help in this case and companies should invest money and time in making the Internet more secure.

To create more awareness and stimulate the companies it should be pointed out to everyone what the risks are that they are currently taking and explaining to them why everyone is dependent on each other. The most ideal place to do this, is to present this during events and meetings, and discuss how network operators collectively should solve this.

Two other things that need to be more thoroughly researched are BGPsec and FlowSpec. BGPsec is currently still in development and the implications are not really clear. FlowSpec could not be tested, due to the lack of test equipment. In future research it could be investigated if there is a lot of overhead with the use of these protocols.

Chapter 12

References

[1] [c-nsp] uRPF Core Internet Routers. http://puck.nether.net/pipermail/cisco-nsp/2013-April/090606.html, April 2013.

[2] Telecommunicatiewet. http://wetten.overheid.nl/BWBR0009950/, 2013. Article 7.4.

[3] ARIN's RPKI Relying agreement. http://mailman.nanog.org/pipermail/nanog/2014-December/071830.html, December 2014.

[4] Albrecht, K. *Dutch Broadband Q3 2014*. Telecompaper, November 2014.

[5] Arends, R., Austein, R., Larson, M., Massey, D., and Rose, S. DNS Security Introduction and Requirements. RFC 4033, IETF, March 2005. https://tools.ietf.org/html/rfc4033.

[6] Baker, F., and Savola, P. Ingress Filtering for Multihomed Networks. BCP 84, IETF, March 2004. https://tools.ietf.org/html/rfc3704.

[7] Callas, J., Donnerhacke, L., Finney, H., and Thayer, R. OpenPGP Message Format. RFC 2440, IETF, November 1998. https://www.ietf.org/rfc/rfc2440.txt.

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

[9] Claise, B., Trammell, B., and Aitken, P. Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information. STD 77, IETF, September 2013. https://tools.ietf.org/html/rfc7011.

[10] Crocker, D., Hansen, T., and Kucherawy, M. DomainKeys Identified Mail (DKIM) Signatures. STD 76, IETF, September 2011. https://tools.ietf.org/html/rfc6376.

[11] Dhamija, R., Tygar, J. D., and Hearst, M. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 581–590.

[12] d'Itri, M. RPSL and rpsltool: Automatic generaction of BGP configurations and filters, 2012.

[13] DMARC. DMARC.org - What is it? http://www.dmarc.org, 2015.

[14] Dmarcian. DMARC Tools for Humans. http://dmarcian.com, 2015.

[15] Dutch Public Prosecution Service. *17-jarige jongen verdacht van hacken KPN*. 2012.

[16] ENISA. Cost of ICT incident: Handbook and Toolset. https://www.enisa.europa.eu/activities/cert/training/training-resources/operational#cost-of-ict-incident, 2015.

[17] ENISA. CSIRT Setting up Guide in English. https://www.enisa.europa.eu/activities/cert/support/guide/files/csirt-setting-up-guide, 2015.

[18] Ferguson, P., and Senie, D. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. BCP 38, IETF, May 2000. https://tools.ietf.org/html/rfc2827.

[19] Galvin, J., Murphy, S., Crocker, S., and Freed, N. Security Multiparts for MIME: Multipart/Signed and Multipart/Encrypted. RFC 1847, IETF, October 1995. https://tools.ietf.org/html/rfc1847.

[20] Goldberg, S. Why is it taking so long to secure internet routing? *Commun. ACM 57*, 10 (2014), 56–63.

[21] Huston, G., and Bush, R. Securing BGP with BGPsec. *The ISP Column* (2011).

[22] Juniper. Configuring Unicast RPF. https://www.juniper.net/documentation/en_US/junos14.2/topics/usage-guidelines/interfaces-configuring-unicast-rpf.html#id-10465509, 2015.

[23] Juniper. Example: Configuring Origin Validation for BGP. http://www.juniper.net/documentation/en_US/junos14.2/topics/topic-map/bgp-origin-as-validation.html, 2015.

[24] Kaspersky. *Global IT Security Risks Survey 2014 Distributed Denial of Service (DDoS) Attacks*. 2015.

[25] Khan, A., Kim, H., and Kwon, T. How Complete and Accurate is the Internet Routing Registry (IRR)? Presentation, Tokyo, Japan, 2011.

[26] Kitterman, S. Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1. RFC 7208, IETF, April 2014. https://tools.ietf.org/html/rfc7208.

[27] KPN. Integrated Annual Report 2013. http://tinyurl.com/p978qml, 2014.

[28] Lepinski, M. BGPsec Protocol Specification. Internet-Draft draft-ietf-sidr-bgpsec-protocol-11, IETF, January 2015. http://www.ietf.org/internet-drafts/draft-ietf-sidr-bgpsec-protocol-11.txt.

[29] Lepinski, M., and Kent, S. An Infrastructure to Support Secure Internet Routing. RFC 6480, IETF, February 2012. https://tools.ietf.org/html/rfc6480.

[30] Lougheed, K., and Rekhter, J. Border Gateway Protocol (BGP). RFC 1105, IETF, June 1989. https://tools.ietf.org/html/rfc1105.

[31] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and McPherson, D. Dissemination of Flow Specification Rules. RFC 5575, IETF, August 2009. https://tools.ietf.org/html/rfc5575.

[32] Miltenburg, W., and Veelenturf, K. Interview with Erik Bais, 2015.

[33] Miltenburg, W., and Veelenturf, K. Interview with Jaya Baloo (KPN CISO department) and Oscar Koeroo (KPN), 2015.

[34] Miltenburg, W., and Veelenturf, K. Interview with Jeroen Veen (KPN Security Architect) and Michel Zoetebier (KPN-CERT), 2015.

[35] Miltenburg, W., and Veelenturf, K. Interview with Marc Gauw, 2015.

[36] Miltenburg, W., and Veelenturf, K. Interview with Rob Vercouteren (KPN-CERT), 2015.

[37] Miltenburg, W., and Veelenturf, K. Interview with Security Engineer 'Mark', 2015.

[38] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and Austein, R. BGP Prefix Origin Validation. RFC 6811, IETF, January 2013. https://tools.ietf.org/html/rfc6811.

[39] NBIP. Beveiliging tegen DDoS aanvallen. http://www.nbip.nl/diensten/nawas-demand-beveiliging-tegen-ddos/, 2015.

[40] NCSC. *Policy for arriving at a practice for Responsible Disclosure*. 2015.

[41] openSPF. openSPF: Implementations. http://www.openspf.org/Implementations, 2015.

[42] openSPF. openSPF: SPF Record Syntax. http://www.openspf.org/SPF_Record_Syntax, 2015.

[43] Prior, M. Configuring routers with RPSL, 2001.

[44] Rekhter, Y., Li, T., and Hares, S. A Border Gateway Protocol 4 (BGP-4). RFC 4271, IETF, January 2006. https://tools.ietf.org/html/rfc4271.

[45] RIPE NCC. Resource Certification (RPKI), Tools and Resources. http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources, 2015.

[46] Riverhead Networks. Whitepaper: Defeating DDoS Attacks. http://www.cse.msu.edu/~cse825/Riverhead_WP.pdf, 2015.

[47] Routing Resilience Manifesto. Mutually Agreed Norms for Routing Security (MANRS). http://www.routingmanifesto.org/wp-content/uploads/sites/14/2014/09/MANRS-PDF.pdf, 2015.

[48] Scarfone, K., and Mell, P. Guide to Intrusion Detection and Prevention Systems (IDPS). *NIST Special Publication 800*, 2007 (2007), 94.

[49] Simon, D. R., Agarwal, S., and Maltz, D. A. AS-based accountability as a cost-effective DDoS defense. *USENIX HotBots* (2007).

[50] Sommer, R., and Feldmann, A. NetFlow: Information loss or win? In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement* (2002), ACM, pp. 173–174.

[51] Stavroulakis, P., and Stamp, M. *Handbook of information and communication security*. Springer, 2010.

[52] Steenbergen, R. Examining the validity of IRR data. NANOG44, Los Angeles, USA, 2008.

[53] Tan, K. M., Killourhy, K. S., and Maxion, R. A. Undermining an anomaly-based Intrusion Detection System using common exploits. In *Recent Advances in Intrusion Detection* (2002), Springer, pp. 54–73.

[54] Terra, F. Responsible Disclosure. http://responsibledisclosure.nl/en/, 2015.

[55] The FENIX Project. Connecting Trusted Networks. http://fenix.zone/en/, 2015.

[56] TIS Labs. Bird BGPsec Implementation. http://bgpsec.tislabs.com/, 2015.

[57] Toonk, A. Securing BGP routing with RPKI and ROAs. http://www.bgpmon.net/securing-bgp-routing-with-rpki-and-roas/, January 2011.

[58] Toonk, A. The Canadian Bitcoin Hijack. http://www.bgpmon.net/the-canadian-bitcoin-hijack/, August 2014.

[59] Toonk, A. Turkey Hijacking IP addresses for popular Global DNS providers. http://www.bgpmon.net/turkey-hijacking-ip-addresses-for-popular-global-dns-providers/, March 2014.

[60] Toonk, A. How accurate are the Internet Route Registries (IRR) — BGPmon, 2015.

[61] Trusted Networks Initiative. Trusted Networks Policy. https://www.thehaguesecuritydelta.com/images/20141124_Trusted_Networks_Policy_beta-vs0_7.pdf, 2014.

[62] Ward, D. Securing BGPv4 using IPsec. Internet-Draft draft-ward-bgp-ipsec-00.txt, IETF, January 2002. https://tools.ietf.org/id/draft-ward-bgp-ipsec-00.txt.

[63] Wikipedia. Sony Pictures Entertainment hack. http://en.wikipedia.org/wiki/Sony_Pictures_Entertainment_hack, 2015.

Appendices

Appendix A

Network Diagram

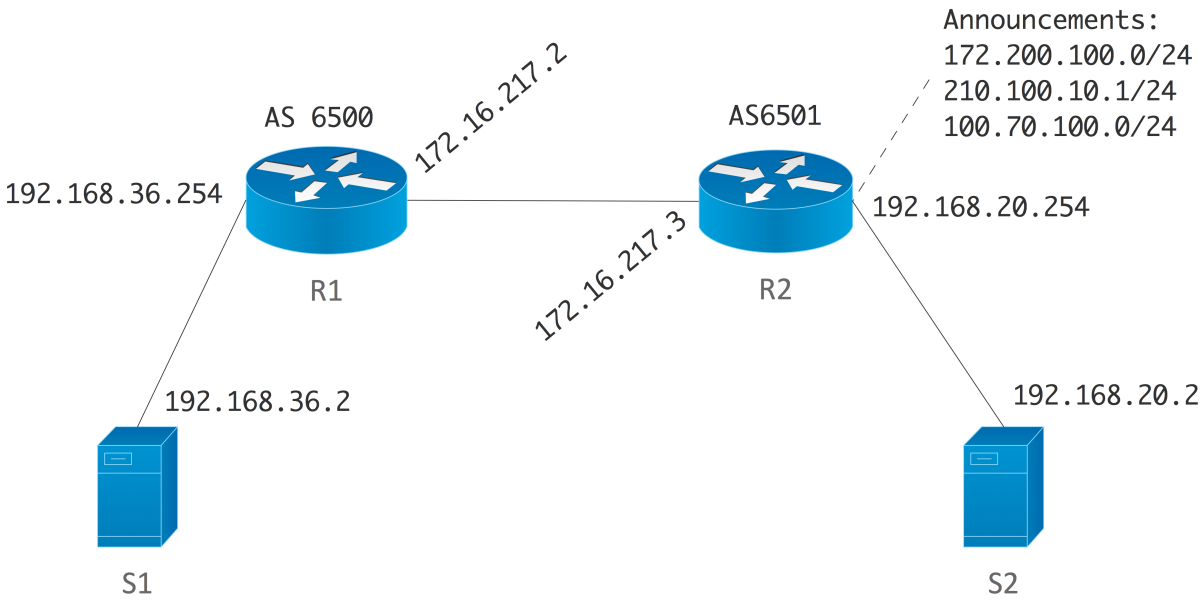

Figure A.1: Example: Network Diagram

Appendix B

Configuration example route filtering

The following config is for a Juniper router and creates filters for the route announcements it receives from its neighbour.

Listing B.1: Juniper ingress and egress filter for route announcements

```
root# set protocols bgp group bgp-peers import import-policy

##The import policy is defined as follows:
root# show policy-options policy-statement import-policy
term 0 {
    from {
        protocol bgp;
        route-filter 192.168.20.0/24 exact;
    }
    then accept;
}
then reject;
```

The above configuration causes the router to only accept the specific *192.168.20.0/24* prefix. If the prefix is more specific, or less like a /22, it would not be accepted by the router. More advanced configuration options exist to allow scenarios where one wants to specify a range that the prefix could have. This example shows how trivial it can be to create an ingress filter. As outlined previously, more advanced configuration options exist. For more information on how to configure the more advanced filters it is advised to take a look at the Juniper TechLibrary.

The same can be done using Cisco, albeit that the configuration is a bit different.

Listing B.2: Cisco ingress and egress filter for route announcements

```
router bgp 6500
##Omitted some config
neighbor 172.16.217.3 distribute-list 101 in

access-list 101 permit ip 192.168.20.0 0.0.0.0 255.255.255.0 0.0.0.0
access-list 101 deny   ip any any
```

The access list specifies that only a *192.168.20.0/24* can be accepted, not less or more specific than a /24. As with Juniper, more advanced configuration exists, where one can specify prefix ranges. For more information on how to configure the more advanced filters it is advised to take a look at the Cisco configuration guides.

Appendix C

Configuration example maximum prefixes

We changed the configuration of our test setup and announced various routes. The upper bound limit was to accept only one route.

Listing C.1: Juniper prefix limit

```
set protocols bgp group bgp-peers family inet unicast prefix-limit
   maximum 1

LOG:
Jan 22 14:32:41    rpd[1147]: 172.16.217.3 (External AS 6501):
   Configured maximum prefix-limit(1) exceeded for inet-unicast nlri: 5
```

The configuration above causes the log entry to show up. With the current configuration one would only get notified that the limit has been reached. It is also possible to terminate the connection, with the *teardown* option, completely when the threshold is met.

Listing C.2: Cisco prefix limit

```
R1(config-router)#neighbor 172.16.217.3 maximum-prefix 1 warning-only
```

The configuration above does the same, but this time for a Cisco router. If one wants to tear down the session when the threshold is met, it is only a matter of leaving the *warning-only* option out of the comment.

Appendix D

Configuration example AS_PATH filtering

In this example we filter all incoming routes and only accept the routes that are originated by AS 6501 and have not been passed through several routers (i.e. there is only one AS in the AS_PATH, which is the neighbour's AS).

Listing D.1: Juniper AS_PATH filtering

```
##Route policy
root# show policy-options policy-statement import-policy
term 0 {
    from {
        protocol bgp;
        as-path 6501-match;
        route-filter 192.168.20.0/24 exact;
        route-filter 210.100.10.0/24 exact;
    }
    then accept;
}
then reject;


##Set the AS-PATH filter
root# set policy-options as-path 6501-match ^6501$
```

Some configuration options have been omitted in the example above, but the route-filters are still in the *import-policy*. Therefore, one can see that it is possible to use different techniques to protect a network from accepting a rogue route. These filters are basically regular expression filters.

The same can be done with Cisco, as shown in the configuration below. In the example below, we have also added the configuration of the second router that prepends an AS number to the prefix *172.200.100.0/24*. Due to the configuration of R1, this route will not be accepted.

Listing D.2: Cisco AS_PATH filtering

```
##R1:
router bgp 6500
 no synchronization
 bgp log-neighbor-changes
 redistribute connected
 neighbor 172.16.217.3 remote-as 6501
 neighbor 172.16.217.3 route-map input-filter in
 neighbor 172.16.217.3 maximum-prefix 1 warning-only
```

```
 no auto-summary
!
ip as-path access-list 10 permit ^6501$
!
ip access-list extended filter
 permit ip host 192.168.20.0 host 255.255.255.0
 permit ip host 172.200.100.0 host 255.255.255.0
 permit ip host 100.70.100.0 host 255.255.255.0
 deny   ip any any
!
route-map input-filter permit 10
 match ip address filter
 match as-path 10



##R2:
router bgp 6501
 bgp log-neighbor-changes
 neighbor 172.16.217.2 remote-as 6500
 !
 address-family ipv4
  redistribute connected
  neighbor 172.16.217.2 activate
  neighbor 172.16.217.2 distribute-list filter in
  neighbor 172.16.217.2 route-map as_prepend out
  no auto-summary
  no synchronization
 exit-address-family
!
ip access-list extended filter
 permit ip host 172.16.217.0 host 255.255.255.0
 permit ip host 192.168.36.0 host 255.255.255.0
 deny   ip any any
access-list 1 permit 172.200.100.0 0.0.0.255
access-list 20 permit any
!
route-map as_prepend permit 10
 match ip address 1
 set metric 100
 set as-path prepend 10
!
route-map as_prepend permit 20
 match ip address 20
```

In the example above some configuration has been omitted (i.e. the interface configuration). Both for Cisco and Juniper more advanced AS_PATH regex matching exists where one can match, for example, if a route has been passed through a certain transit. For more information it is advised to take a look at the configuration guides of the according vendor.

Appendix E

Configuration example MD5 and IPsec

The example below, shows how MD5 authentication can be configured for BGP peers.

Listing E.1: Cisco MD5 authentication

```
R1(config)#router bgp 6500
R1(config-router)#neighbor 172.16.217.3 password example

##One needs to do this on both routers
R2(config)#router bgp 6500
R2(config-router)#neighbor 172.16.217.2 password example

##Error:
*Mar  1 01:48:25.887: %TCP-6-BADAUTH: No MD5 digest from
   172.16.217.3(179) to 172.16.217.2(29172)
```

Please note to not copy paste this configuration, as the password is too simple and will thus be the same as this document. The configuration is trivial and one can expect an error in the logs, as shown above, when the peer has not configured the password yet.

The same can be done of course with Juniper, as shown below.

Listing E.2: Juniper MD5 authentication

```
##Need to do the following action on both routers
[edit protocols bgp group bgp-peers]
root# set authentication-key example

##Error:
Jan 27 16:45:24    /kernel: tcp_auth_ok: Packet from 172.16.217.3:65164
   unexpectedly has MD5 digest
```

For completeness, the error message that one could expect is shown as well.

As discussed earlier, it might be better to use IPsec for authentication and confidentiality. In the example below we configure an IPsec in transport mode to encrypt the BGP traffic.

Listing E.3: Juniper IPsec protection for BGP

```
##One needs to do this on both peers

protocols {
    bgp {
        group bgp-peers {
            ipsec-sa protect-bgp;
```

```
                        ##Omitted output
        }
    }
}

security {
    ipsec {
        security-association protect-bgp {
            mode transport;
            manual {
                direction bidirectional {
                    protocol esp;
                    spi 1000;
                    encryption {
                        algorithm 3des-cbc;
                        key ascii-text
                            "$9$rBKKWxbs4Di.Ndi.P56/lKML7VgoGqmTwYmTz3t
                            pWLx-b2ZUH5Qn4aQn/CB17-Vw4aGDi.fT"; ##
                            SECRET-DATA
                    }
                }
            }
        }
    }
}


##Error one can get:
Jan 27 21:20:23  R2 /kernel: IPv4 ESP input: no key association found
   for packet(SPI=1000 seq=49 src=172.16.217.2 dst=172.16.217.3)
```

As the comment outlined, both the configuration mentioned above, needs to be configured on one's router and the peer's router. Incorrectly configuring this, can result in not having a BGP session or error messages in the log. The log above shows that there has been a configuration error. In this case it was a mistyped IP address, but as always, consult the logs if it is not possible to create a session between the two peers.

The same can be done on a Cisco device, as shown below.

Listing E.4: Cisco IPsec protection for BGP

```
##R1 config:
!
crypto isakmp policy 10
 authentication pre-share
crypto isakmp key 6 ciscociscocisco123 address 172.16.217.3
!
!
```

```
crypto ipsec transform-set esp-des esp-3des
 mode transport
!
crypto map protect-bgp 10 ipsec-isakmp
 set peer 172.16.217.3
 set transform-set esp-des
 match address ipsec_secured
!
interface FastEthernet0/1
 ip address 172.16.217.2 255.255.255.0
 duplex auto
 speed auto
 crypto map protect-bgp
!
ip access-list extended ipsec_secured
 permit tcp host 172.16.217.2 host 172.16.217.3
!

##R2 config:
!
crypto isakmp policy 10
 authentication pre-share
crypto isakmp key 6 ciscociscocisco123 address 172.16.217.3
!
!
crypto ipsec transform-set esp-des esp-3des
 mode transport
!
crypto map protect-bgp 10 ipsec-isakmp
 set peer 172.16.217.3
 set transform-set esp-des
 match address ipsec_secured
!
interface FastEthernet0/1
 ip address 172.16.217.2 255.255.255.0
 duplex auto
 speed auto
 crypto map protect-bgp
!
ip access-list extended ipsec_secured
 permit tcp host 172.16.217.2 host 172.16.217.3
!
```

The example above shows the different configurations of the two peers. It is a bit more work and in fact all TCP traffic that is initiated from a router, so not routed traffic will be send over the IPsec tunnel.

Appendix F

Configuration example BGP Origin Validation

We made small changes to the network, in order to provide the reader with a configuration example, as the second router now announces the route *181.50.0.0/22* and has AS number *10620* instead of AS number *6501*. Moreover, the server connected to the first router runs the RPKI validator service created by RIPE NCC [45].

The configuration below shows how one can configure their router to create an import filter for routes that are announced by its peer.

Listing F.1: Cisco BGP Origin Validation

```
router bgp 6500
 bgp log-neighbor-changes
 bgp rpki server tcp 192.168.36.2 port 8282 refresh 600
 neighbor 172.16.217.3 remote-as 10620
 neighbor 172.16.217.3 route-map rpki-loc-pref in
!
route-map rpki-loc-pref permit 10
 match rpki invalid
 set local-preference 90
!
route-map rpki-loc-pref permit 20
 match rpki not-found
 set local-preference 100
!
route-map rpki-loc-pref permit 30
 match rpki valid
 set local-preference 110
!
```

The *bgp rpki server* command states the server that is going to be used by the router for validating the routes. The route maps define what needs to be done with a certain route, in this case setting the local preference. More advanced configurations can be created where one can, for example, deny routes that are considered *invalid* by the validator.

After reconfiguration, the second router now announces *181.50.0.0/22*. As one can see below, the route has been validated with RPKI and its state is considered *valid*.

Listing F.2: Cisco Route Validated

```
R3#show ip bgp
BGP routing table entry for 181.50.0.0/22, version 36
Paths: (1 available, best #1, table default)
```

```
  Not advertised to any peer
  Refresh Epoch 1
  10620
    172.16.217.3 from 172.16.217.3 (192.168.200.1)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      path 67A3A460 RPKI State valid
      rx pathid: 0, tx pathid: 0x0
```

The test equipment that we had access to, did not support BGP Origin Validation. Therefore, no configuration example is provided. However, if one wants to configure Origin Validation on a Juniper router it is advised to consult the following Juniper TechLibrary [23].

Appendix G

Configuration example uRPF

The following example, in Listing G.1, contains two possible configurations for uRPF on Juniper routers for a specific interface. The first one, is for uRPF strict-mode, where the second one is for uRPF loose mode.

Listing G.1: uRPF Junos example

```
admin# set interfaces ge-0/0/0 unit 0 family inet rpf-check

admin# set interfaces ge-0/0/0 unit 0 family inet rpf-check mode loose
   #For loose mode
```

The following example, in Listing G.2, contains two possible configuration for uRPF on Cisco routers for a specific interface. The first one, is for URPF strict-mode, where the second one is for uRPF loose mode.

Listing G.2: uRPF Cisco example

```
R1(config)#interface FastEthernet0/1
R1(config-if)#ip verify unicast source reachable-via rx

##Loose mode:
R1(config-if)#ip verify unicast source reachable-via any
```

Appendix H

Configuration example SPF

Since there are a lot of different mail servers available, there are also a lot of different SPF implementations available. openSPF published a list with SPF extensions for a couple of mail servers [41]. Besides the software, the DNS record needs to be created according to the SPF standard in RFC 7208 [26].

The DNS record consists of a few elements. Every SPF records starts with the version of SPF, namely *v=spf1*. The second part of the SPF record specifies the IP address that is allowed to send email on behalf of the domain. It is also possible to configure a specific IP range using a prefix-length, for instance *85.12.6.41/30*. The last part of the DNS record specifies the qualifier. In the example in Listing H.1, only the server with IPv4 address *85.12.6.41* and IPv6 address *2a01:788:f009::85:12:6:41* is allowed to send mail for the *exampledomain.com* domain. All other servers are prohibited from sending email on behalf of the *exampledomain.com* domain. Besides these simple options, there are more ways to configure the SPF DNS record. These options can be found in RFC 7208 [26].

Listing H.1: SPF DNS record example

```
"v=spf1 ip4:85.12.6.41 ip6:2a01:788:f009::85:12:6:41 ~all"
```

Appendix I

Configuration example DKIM

Just like with SPF, there are a lot of different implementations possible for DKIM. It highly depends on the kind of mail server. A common implementation for DKIM is openDKIM. First it is important to create the DKIM public and private key. To use the key generator, one needs to install *opendkim-tools*.

Listing I.1: DKIM key generation

```
cd /etc/opendkim/keys/exampledomain.com/
opendkim-genkey -D /etc/opendkim/keys/exampledomain.com/ -d
    exampledomain.com -s default
```

As shown in Listing I.1, a key will be generated for *exampledomain.com*. The '-D' flag specifies the directory in which the keys need to be saved, the '-d' flag specifies the domain name, and the '-s' flag specifies the name of the key, which also needs to be used for creating the TXT record. In case of the example, the TXT record label would be *default._domainkey.exampledomain.com*, as shown in Listing I.2.

Listing I.2: DKIM DNS record

```
default._domainkey IN TXT "v=DKIM1; k=rsa;
    p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDGCMlMi5hSxFlpb3FlWElQlTr
    HziSh4kzELQyIAp92+nEP2iXNURCqRFJEeoOZ0Qj2uM3wr7NOLabxD+o0CjKtd0u5e
    RU2YWD/dQln6lY+Lt8N+Du+haVgzGzoj/XGKsBUJ2aRc2JnmBxecE0b/3o3n4doobrN
    e5se/K4cINA31QIDAQAB" ; ----- DKIM key default for exampledomain.com
```

The key generation will create a file named *default.private* and a file named *default.txt*. The *default.private* file is the private key for the mail server, and the *default.txt* contains the DNS TXT record, as shown in Listing I.2. This record contains the public key for the mail server in the field *p*. After creating the keys, the configuration file *opendkim.conf* should be altered to the proper settings, and the *KeyTable* needs to contain the DNS record label and the location of the public key.

The last step is to configure the mail server to verify the DKIM signatures. This depends on the kind of mail server that is running. Since configuring DKIM depends on the kind of mail server, it is advised to check the documentation for that specific mail server.

Appendix J

Configuration example source filtering

The configuration shown below is made for router 1. Only packets that can be seen as 'input/ingress' to interface FastEthernet0/0 and have a source IP address in the range of *192.168.36.0/24*, will be accepted.

Listing J.1: Cisco source filtering

```
##Some configuration has been omitted
!
interface FastEthernet0/0
 ip address 192.168.36.254 255.255.255.0
 ip access-group block_spoof in
!
ip access-list extended block_spoof
 permit ip 192.168.36.0 0.0.0.255 any
 deny   ip any any
!
```

The same is possible with Juniper. Routes that can be seen as 'input/ingress' to interface ge/0/0/0 and have a source IP address in the range of *192.168.36.0/24*, will be accepted.

Listing J.2: Juniper source filtering

```
##Some configuration has been omitted
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                filter {
                    input source-filter;
                }
                address 192.168.36.254/24;
            }
        }
    }
}

firewall {
    family inet {
        filter source-filter {
            term 0 {
                from {
                    source-address {
```

```
                        192.168.36.0/24;
                    }
                }
                then accept;
            }
            term default {
                then {
                    reject;
                }
            }
        }
    }
}
```

For more advanced configurations it is advised to take a look at the configuration guides of the according vendor.