Uncovering the vulnerable

Exploring the issue of TCP reflective amplification in the network of an ISP

J. Oortwijn

TU Delft

Uncovering the vulnerable

Exploring the issue of TCP reflective amplification in the network of an ISP

by

J. Oortwijn

Master thesis submitted to Delft University of Technology

in partial fulfilment of the requirements for the degree of

Master of Science

in **Engineering and Policy Analysis**

to be defended publicly on April 17, 2023.

Student number:	4593472	
Thesis committee:	Dr. ir. C. Hernandez Gañán,	TU Delft, Chair
	Dr. R. S. van Wegberg,	TU Delft, First supervisor
	Dr. ir. J. De Stefani,	TU Delft, Second supervisor
	F. G. A. Groenewegen, MSc.,	KPN, External supervisor

TU Delft

kpn

Preface

This thesis serves as my conclusion to the master of Engineering and Policy Analysis at the Delft University of Technology. Over the past months, I have dived into the topic of cyber security and the role of internet service providers in this complex environment. An interesting journey which certainly taught me a lot which I'm thankful for. This would not have been possible without the help of some people who I would like to thank.

First of all, I would like to express my gratitude to my supervisors for their guidance and support during this process. Foremost, I am deeply grateful to Carlos Hernandez Gañán for introducing me to this topic and providing me with invaluable guidance and support throughout the entire process. I really appreciate the fact that you were always available for ad-hoc meetings to discuss my progress and steer me in the right direction with helpful comments and suggestions. I also would like to thank Rolf van Wegberg and Jacopo De Stefani for serving as first and second supervisors. Even though we did not have frequent meetings, I believe that your feedback improved my work a lot by highlighting topics and points of improvement from another perspective.

Moreover, I would like to thank Anand Groenewegen and Dennis van Drunen of the KPN abuse team for offering me the opportunity to conduct this project at KPN. The feedback provided during our weekly update meetings, as well as the answers to all my questions, have been valuable and extremely helpful. Your (technical) assistance and insights greatly contributed to my thesis. Moreover, I would like to thank the abuse team for always being available for my questions and for your help throughout the project, which above all, made my time as an intern an enjoyable and memorable experience.

Lastly, I would to thank all my friends and family who patiently listened to my discussions (and complaints) about my thesis even though there were times you probably had no clue what I was talking about. Your support helped me to release some steam and gain perspective, and I am truly grateful.

J. Oortwijn
Rotterdam, April 2023

Summary

The rapid growth of internet-connected devices has led to a significant increase in the number of cyber attacks, resulting in security challenges related to IoT. Researchers have discovered a new attack technique that can be used for launching large DDoS attacks, which involves TCP reflective amplification attacks by abusing middleboxes and IoT devices. This newly discovered attack technique offers a novel perspective on the abuse of TCP, potentially surpassing the effectiveness of UDP-based amplifiers in terms of amplification factors. This highlights the potential for a significant impact of this technique. In order to assist Internet Service Providers (ISP's) in mitigating this vulnerability present at their customers, a deeper understanding of this novel attack technique is needed. Since this attack technique was recently discovered, limited information is available about the vulnerability itself, the end-users who operate these vulnerable devices, and the characteristics of both the users and the devices that are contributing to the problem.

The thesis primarily focuses on exploring vulnerable devices and their end-users within the consumer network of a Dutch ISP, KPN. The ultimate goal is to gather more information on the types of vulnerable devices and actors involved to eventually assist an ISP in making informed decisions to remediate the vulnerability in their network. A mixed-method approach is used to analyse both the quantitative and qualitative data available. First, quantitative data on vulnerable devices in the ISP network is analysed to identify customer IP addresses, track them over time, and perform network scans to characterise the vulnerable devices. Second, end-users updating behaviour, notification perception, and characteristics are analysed through a notification campaign and interviews with KPN customers and a manufacturer. This data is used to answer the research questions: *To what extent can we characterise the issue of TCP reflective amplification in the network of an ISP?*

The study found that the problem can be described in two different issues: vulnerable middleboxes and vulnerable consumer IoT devices with broken TCP protocols. The problem of vulnerable middleboxes has been solved in the network of the Dutch ISP as manufacturers have released updates remediating the vulnerability. This is not the case for vulnerable consumer IoT, as updating consumer IoT devices does not necessarily address the vulnerability present in the devices that have been identified. However, vulnerability notifications can potentially be useful for end-users to encourage them to update their vulnerable devices. The analysis based on the COM-B model showed that the interviewees possess both the capability and motivation to execute the steps outlined in a vulnerability notification. Nevertheless, it was found that vulnerability notifications should include a comprehensive description of the vulnerable device and the specific vulnerability, enabling end-users to accurately identify the device in question and perform the necessary software or firmware update.

The study highlights the presence of vulnerable devices in the ISP network that cannot be remediated by updating the device due to the unavailability of a fix. This calls for the exploration of alternative notification methods like walled garden notifications for ISP's to address the issue as mail notifications seem not feasible at the moment of writing. While updating devices is a suggested solution, it may not be feasible for end-users with vulnerable consumer IoT devices, making it crucial for manufacturers to ensure their products have secure TCP protocols. While end-users are motivated and capable to keep their vulnerable devices up to date, whether or not they receive a vulnerability notification from their ISP, this action alone will not fully address the vulnerability as long as manufacturers remain unaware of the issue or fail to provide updates to remedy it.

Contents

List of abbreviations

(D)DoS	(Distributed) denial of service
AS	Autonomous system
BMS	Building management system
BM	Business market
CM	Consumer market
COM-B	Capability, opportunity, motivation and behaviour
DoS	Denial of service
HTML	Hypertext markup language
HTTP	Hypertext transfer protocol
IDS	Intrusion detection system
IoT	Internet of things
IP	Internet protocol
ISP	Internet service provider
TCP	Transmission control protocol
UDP	User datagram protocol
URL	Uniform resource locator
WM	Wholesale market

List of figures

List of tables

1 Introduction

1.1. Research background

The amount of devices connected to the internet has been rapidly growing and will continue to do so. In 2022, it is estimated that there are around 13 billion internet-connected devices in the world and this is expected to grow up to 30 billion in 2030 (Transforma Insights, 2022). Although these devices aim to improve our lives in various ways, the rise of internet-connected devices also increased the attack surface for cyber criminals to exploit (Anand et al., 2020). Over the past years, the number of attacks on internet-connected devices has grown substantially with a jump of approximately 900 percent of such attacks in 2019 (Prajapati et al., 2022).

There are numerous security challenges related to internet-connected devices or the Internet of Things (IoT) which are extensively studied by academics. These challenges are often described in terms of security issues and privacy issues with IoT (Maple, 2017; Omolara et al., 2022; Panahi Rizi et al., 2022). Security challenges for IoT are present because of three main reasons (Atzori et al., 2010). First, IoT items are unattended most of the time. Second, because of their nature, most IoT devices have low capability in terms of energy usage and computation power. For this reason, they cannot implement security measures such as data encryption and use lightweight insecure protocols. Lastly, communication between IoT happens mainly wireless over a large network of different interconnected devices, generating a large attack surface (Rizvi et al., 2020).

When looking at the past, there have been countless examples of security breaches of internet-connected devices resulting in large societal disruptions. In 2016, DNS provider Dyn was hit by a large DDoS attack using a botnet of many IoT devices, resulting in the inaccessibility of numerous popular websites such as Netflix, Twitter and PayPal (Sinanovic et al., 2017). Denial of service (DoS) attacks are one of the most common cyber attacks utilised today (Y. Li et al., 2021). DoS attacks attempt to make a device or the network resources unavailable to its legitimate clients by temporarily or permanently disrupting the services of a host connected to the internet by flooding the victim with traffic (Kaur Chahal et al., 2019; Nosyk et al., 2022). Distributed denial of service (DDoS) attacks use the same principle but for this type, multiple hosts (e.g. via a botnet) are used to flood a victim with traffic instead of one single source (Douligeris et al., 2004).

On a daily basis, newly discovered security vulnerabilities and novel forms of attacks are detected. Bock et al. (2021) propose in their article a new attack technique that, by abusing middleboxes, can be used for launching large DDoS attacks. This paper explains how these middleboxes can be used for TCP reflective amplification attacks. A middlebox is defined as an intermediary device that performs functions beyond the typical, standard operations of an IP router on the datagram path between a source host and a destination host (Carpenter et al., 2002). What makes this type of DDoS attack different from attacks currently seen in the wild is the way how this attack is executed. Normally UDP based protocols are abused for creating an amplification of data which is sent to the victim, while the attack described by Bock et al. (2021) is TCP based. There is a common consensus that UDP based amplifiers are superior compared to TCP based amplifiers when looking at the amount of amplification

that can be achieved (Kührer et al., 2014a; Kührer et al., 2014b; Nosyk et al., 2022). Bock et al. (2021) explains that, by using middleboxes, the achieved amplification factors can be similar to UDP based amplifiers or even higher. This discovery poses a serious security threat since using TCP for transmitting data is more common than UDP and thus more options for attackers are available to send DDoS attacks (Nguyen et al., 2018).

While the paper by Bock et al. (2021) focuses on middleboxes as a tool for exploiting this new attack technique, any internet-connected device with a broken TCP implementation could be abused by malicious people for launching TCP reflective amplification attacks. Academics already studied IoT devices in households that can act as possible amplifiers for launching DDoS attacks using TCP protocols. For example, Lyu et al. (2017) show that consumer IoT can be used for TCP amplification attacks but conclude that attacks using UDP are more effective. Furthermore, the papers by Kührer et al. (2014b) and Kührer et al. (2014a) show that TCP protocols can be exploited for reflective amplification attacks. The novel attack technique described by Bock et al. (2021) provides a completely new perspective on abusing services on TCP as they show that TCP reflective amplification attacks can be significant in size and therefore at least as interesting as UDP-based amplifiers for launching DDoS attacks.

Three key stakeholders can be identified who are involved in the problem of IoT devices which can be abused for TCP reflective amplification attacks. First, the manufacturers of these devices as the problem is created by a poor configuration of the device itself. Second, end-users are responsible for all the changes they implement to their devices such as software updating and port forwarding. Lastly, internet service providers (ISP's) detected security issues on their network and are tasked with mitigating vulnerabilities that could harm their customers, their network or both.

Internet service providers can play an important role in helping people with vulnerabilities in their internet network due to their unique position in the market (Fruchter, 2019; Moore et al., 2009; Rowe et al., 2011). ISP's are able to link identified vulnerable IP addresses to contact information of actual people and use their different communication channels to notify the customer about the vulnerability. Notifying people about the vulnerability of their devices has been a suggested method to mitigate the vulnerability (Cetin, Ganan, Korczynski, et al., 2017; Hennig et al., 2022; F. Li, Bailey, et al., 2016). An ISP can take various measures to encourage customers to mitigate the vulnerability in their network. For example, an ISP could send an email notification, a direct message or notify the customer of interest via quarantining their network (Cetin, Ganan, Altena, Tajalizadehkhoob, et al., 2019; F. Li, Bailey, et al., 2016). The different notification measures, and the content of the notification, require the ISP to make a well-balanced decision in order to decide which mechanism to adopt and how to formulate the content of such vulnerability notification.

Before notifying a customer about a potential security issue, the issue itself should be thoroughly understood. This helps the ISP to make this well-balanced decision on what steps to take in order to remediate a vulnerability of interest. In the case of vulnerable middleboxes or other IoT devices with broken TCP protocols, this information is at an undesired level potentially obstructing ISP's from effectively taking action to solve the issue.

1.2. Problem statement

As described by Bock et al. (2021), TCP reflective amplification (by abusing middleboxes) can lead to amplification factors that surpass those typically observed with commonly used TCP or UDP-based amplifiers. Nevertheless, little is known about this type of DDoS attack. At the moment of writing, there are not many known cases of these attacks being used in the wild. Still, it is expected that malicious people are currently adopting and fine-tuning this attack technique (CIRT Akamai, 2022). Therefore, intervention is needed to prevent large DDoS attacks from happening. ISP's are actors in this problem who can address this issue due to their unique position on the internet. However, there are problems that restrain ISP's from effectively helping their customers with remediating the vulnerability. The main cause for these issues is the lack of information on this vulnerability.

There are unknowns about the characteristics of the vulnerable devices and the vulnerability itself. The vulnerable device could be anything ranging from a middlebox to a consumer IoT device. Furthermore, in the case of middleboxes, these devices are located somewhere between a source host and a destination host, thus one can not simply scan the host IP to identify and characterise the vulnerable device.

This issue is already described in previous research. Nosyk et al. (2022) found routing loops that can be used for sending a packet flood towards a victim, exploiting the above-mentioned attack technique. They explain that more research is needed to identify and characterise the devices which cause the problem. Moreover, despite that TCP reflective amplification attacks are often used for launching DDoS attacks, which involve sending TCP packets to a vulnerable server or device that reflects the data by overloading the victim with a large number of SYN-ACK packets, there exists a scarcity of research in this area (Kührer et al., 2014b; Mohana Priya et al., 2014).

Because of this lack of information, an ISP can not effectively communicate the vulnerability to its customers with specific steps to solve the issue. For vulnerability notifications to work, ISP's should have the right information to communicate, manufacturers should have released a security patch and end-users should perform certain behaviour to secure their vulnerable devices. Therefore, it is interesting to research these three actors and their roles in regard to the problem of TCP reflective amplification. These knowledge gaps are further addressed in chapter 1.5.

1.3. Research objective and questions

This research tries to address knowledge gaps which are linked to internet-connected devices vulnerable to TCP reflective amplification. As explained in chapter 1.2, TCP reflective amplification is not widely covered by academics. This research tries to address this lack of information by analysing the problem of devices which can be abused for launching TCP reflective amplification attacks in the consumer network of an ISP. The main objective of this thesis is to explore the issue of vulnerable middleboxes and IoT devices which can be abused for this type of attack in terms of vulnerable devices characteristics and end-users behaviour in the network of an ISP. The goal is to gain more information on the type of vulnerable devices which can be abused for the attack technique of interest, the different key actors involved and their role in the problem to ultimately help an ISP with making well-informed decisions on how to act in order to remediate the vulnerability in their network. This objective can be divided into the following four sub-objectives:

- The first sub-objective is to characterise the vulnerable devices in the consumer network of an ISP. As stated before, little is known about these devices such as the type or the brand. By doing so, information is gathered which could help an ISP with notifying end-users to communicate specific remediation steps.
- The second sub-objective is to map the characteristics of the end-users who own a vulnerable middlebox or IoT device. This will provide information on what type of end-user operate these vulnerable devices which may help ISP's with making an informed decision on how to act.
- The third sub-objective is to identify the security behaviour of these end-users, and more specifically their patching behaviour in regard to their vulnerable device. Patching is a solution suggested, and proven to be effective when remediating vulnerable middleboxes (Pal, 2022). As this vulnerability is the result of a broken TCP protocol, patching is assumed to be a solution for fixing IoT devices, other than middleboxes, which can be abused for launching TCP reflective amplification attacks.
- The fourth, and last, sub-objective is to gain an understanding of what information end-users need to be able to remediate the vulnerability. This will provide background information for possible future vulnerability notifications which ISP's can send to inform end-users. This is interesting to understand as notifying end-users is the main instrument for ISP's to address the issue in their network.

Based on this objective and the sub-objectives the following research question is formulated which this research aims to answer:

To what extent can we characterise the issue of TCP reflective amplification in the network of an ISP?

To answer the aforementioned main research question, the following sub-questions have been formulated. Below each sub-question is presented and briefly discussed.

SQ1: *What does the population of vulnerable devices look like in the consumer network of an ISP?*
To start off, it is necessary to examine the characteristics of the vulnerable devices with respect to their

device types to identify which devices are responsible for the observed amplification. For this sub-question, previously identified IP addresses are included that are considered remediated at the time of writing. It is valuable to understand why certain devices have been fixed for this vulnerability while others remain vulnerable. This contextual information can aid an ISP in developing content for a future vulnerability notification.

SQ2: *What are the characteristics of vulnerable device end-users?*
The second-sub question tries to identify the characteristics of vulnerable device end-users. This includes age, gender, and perceived skills in computer security. This can give an understanding of what type of consumer operates these vulnerable devices. Furthermore, it is interesting to understand for what purposes (business or private), end-users use the vulnerable devices. Especially for middleboxes, it is not expected that normal consumers use these more complex devices for private purposes. Similar to the first sub-question, previously identified vulnerable IP addresses, and thus customers, are included that are considered remediated when this research took place.

SQ3: *What does the updating behaviour of vulnerable device end-users in the consumer network of an ISP look like?*
The suggested resolution for addressing the vulnerability in middleboxes is to apply patches. Subsequently, an ISP may recommend this solution to its customers who own a vulnerable IoT device or middlebox. Thus, it is important to examine the updating behaviour of end-users, including whether they keep their (smart) devices up to date and what motivates them to perform or neglect updates.

SQ4: *What is the effect of a vulnerability notification on vulnerable device end-users in terms of performed actions and perception?*
In order to assess the possible impact of a vulnerability notification on end-users of vulnerable devices, it is crucial to comprehend their perception of such notifications and their response to them. This approach can aid in identifying the requirements of end-users for identifying vulnerable devices and determining the information they need to respond appropriately to the notification received. These insights can shed light on how a potential vulnerability notification would affect end-users and whether it is practical for an ISP to implement this mechanism to remedy the vulnerability in their network.

1.4. Research approach

The objective of this research is exploratory in nature, therefore both the available quantitative and qualitative data are taken into account when answering the research questions. A mixed-method approach allows this study to combine qualitative and quantitative methods within different phases of the research process (Terrell, 2012). Furthermore, these methods are applied in the context and network of KPN, a Dutch ISP. Focusing on this single case allows this research to study the issue of internet-connected devices, vulnerable to being abused for TCP reflective amplification attacks, extensively. The sequential transformative strategy is applied as first the quantitative data on the vulnerable devices is analysed and sequentially, qualitative data is collected and analysed based on the findings in the first quantitative phase (Creswell, 2014). For this research, the emphasis lies on the second phase, the qualitative data. The research can be divided into four phases in which different methods are used to analyse either the qualitative or quantitative data available. Below each of the four phases is briefly discussed.

- First a literature study is performed to gain an understanding of the context of this research. Academic literature on the security of internet-connected devices, TCP reflective amplification attacks, vulnerability notifications and end-user security behaviour will provide the theoretical background on which this research is founded. Furthermore, this will offer a thorough understanding of the issue that is essential for creating vulnerability notifications that will be sent to end-users in a subsequent phase of the research.
- Second, quantitative data provided on vulnerable devices is analysed to identify the IP addresses in the network of KPN, track these over time and analyse these devices. This data is complemented by data retrieved by performing network scans on the IP addresses of interest. Together this data will be used to characterise the vulnerable devices.
- The third phase involves the analysis of end-users updating behaviour, notification perception,

and characteristics to assess their role. During the course of 5 weeks, all customers with vulnerable devices will receive notifications regarding the vulnerability, while end-users of previously vulnerable devices will also be reached out to. Subsequently, each contacted KPN customer is invited to participate in an interview. The qualitative data obtained from these interviews will provide insights into the updating behaviour of the interviewees. Furthermore, their perception of the received notification and their actions can be analysed to offer recommendations for the feasibility of implementing vulnerability notifications for ISP's.

- In the fourth phase, all the data obtained during the preceding three phases will be analysed and processed. The findings derived from this analysis will serve as the foundation for addressing the research question and sub-questions.

Chapter 4 presents the complete methodology of this research including a comprehensive view of the four above-mentioned phases. The complete overview of this research, and the different phases, is visualised below in figure 1.

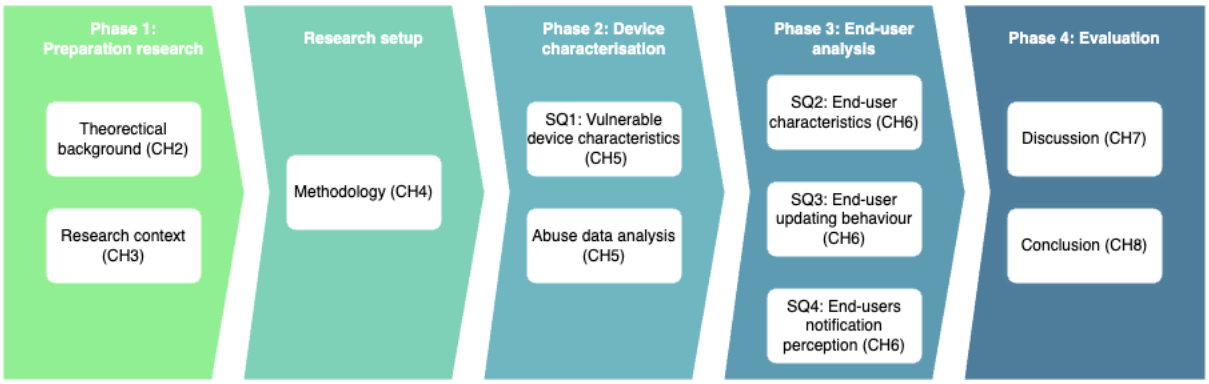

Figure 1: Research flow diagram

1.5. Research relevance

This subsection discusses the importance of this research in terms of its potential contributions to scientific and societal knowledge. Additionally, the benefits of this research for KPN, the Dutch ISP that facilitated the study, as well as other ISP's, are emphasised.

1.5.1. Scientific relevance

As discussed before in chapter 1.2, there is little literature available discussing the issue of devices vulnerable to TCP reflective amplification attacks. This research tries to address this knowledge gap by analysing this topic in combination with looking at the role of an ISP and end-users. At the moment of writing no previous research has been conducted into the research objective as described in chapter 1.3.

In the past, research has been conducted into vulnerable middleboxes or IoT devices and, the attack related to these vulnerable devices, TCP reflective amplification attacks (Bock et al., 2021; Islam et al., 2022; Kührer et al., 2014a; Kührer et al., 2014b; Nosyk et al., 2022). Nevertheless, none of these studies analysed the problem of these devices in combination with the role of an ISP. Moreover, previous studies that focus on TCP reflective amplification attacks, such as Kührer et al. (2014a), Kührer et al. (2014b), and Mohana Priya et al. (2014), primarily discuss SYN-ACK floods, whereas this research centres on broken TCP protocols that respond similarly to vulnerable middleboxes. This topic will be further elaborated upon in chapter 2.

Furthermore, there have been studies into vulnerability notifications by ISP's to help end-users with a vulnerability or malware cleanup before but not in the context of the vulnerability discussed in this thesis (Cetin, Ganan, Altena, Kasama, et al., 2019; Rodríguez et al., 2021; J. Zhang et al., 2017). By exploring the issue within an ISP, this thesis tries to provide background information which ISP's can use for potential future vulnerability notifications. Hence, this can be regarded as the initial step in

mitigating the vulnerability. The subsequent step would be to assess the effectiveness of vulnerability notifications which is out of the scope of this thesis.

1.5.2. Societal relevance

Enhancing the cyber security of vulnerable devices and their end-users is not only in favour of an ISP but a wider range of stakeholders. While an ISP can easily notify the end-users about the vulnerability, the exploitation of the vulnerability does not affect the ISP or end-user directly. A DDoS attack, leveraging these vulnerable devices, could harm any organisation or individual in their lives by obstructing digitised services such as banking or shopping. As a DDoS attack can be launched against anything or anyone the impact of vulnerable devices, which can be used for such attacks, is practically boundless. A better understanding of this attack can help ISP's to better inform their customers and reduce the number of vulnerable devices in their network thus contributing to mitigating cybercrime. As an ISP can be seen as a line in the defence against such attacks, any potential victim would benefit from this research.

This research addresses an issue which is related to the broader topic of cyber security. Cyber security is perceived as critical to both people's prosperity and security. The importance of cyber security is stressed by multiple large public organisations such as UN (2023) and European Commission (2017). Furthermore, this thesis tries to identify the patching behaviour of vulnerable device end-users. Last year, the Dutch national government relaunched an awareness campaign to stimulate people to update their smart devices as it was found that more than half of the Dutch population lacks certain patching behaviour (Rijksoverheid, 2021).

This research aims to improve the understanding of vulnerable devices that can be used for TCP reflective amplification attacks in the network of an ISP. The findings of this research can assist the ISP and potentially other ISP's to effectively mitigate the risks associated with this vulnerability. Doing so makes a small but important step towards a safer cyber world.

1.5.3. Added value for KPN

The gain in knowledge about vulnerable devices will help KPN with better addressing the issue and helping their customers with these devices. The responsibility of managing security incidents in their network has been delegated to the KPN abuse team. More information on these security issues can subsequently help them with making better-informed decisions on how to respond. The tools used for performing the analysis are provided by KPN and thus recommendations based on the results could be directly implemented in the working procedures of the KPN abuse team. Improved working of the KPN abuse team, could result in a gain in time efficiency and cost efficiency as KPN abuse staff members need less time to help customers with their vulnerable devices.

Additionally, this research could help other IPS's with devices in their network vulnerable to being abused for launching TCP reflective amplification attacks. Information on the vulnerability in general and customer reactions and attitudes towards vulnerability notifications can help any other IPS's with improving their own services.

1.6. Link to EPA programme

This thesis is written as part of the requirements for the degree of Master of Science in Engineering and Policy Analysis (EPA) at the Delft University of Technology. This research fits the EPA programme because it addresses the grand challenge of cyber security. The focus of this thesis is on a vulnerability that exists within middleboxes and IoT devices which is of a technical nature and arises from the misconfiguration of these devices' settings, potentially leading to severe cyber attacks. Integrating the technical and social aspects of this vulnerability is crucial as these cyber attacks can have a significant impact on individuals, corporations, and society as a whole. This thesis aims to analyse the TCP reflective amplification vulnerability and provide policy recommendations to ISPs for addressing this issue. Furthermore, the vulnerability can be mitigated by people performing certain behaviour thus addressing the societal aspect of it. The EPA program teaches several tools to analyse this complex socio-technical system. In line with the program, this research combines quantitative data analysis with a qualitative interview method.

1.7. Thesis outline

The remainder of this thesis is as follows. Chapter 2 provides the theoretical background for this study and discusses previous academic work related to the topic, in order to contextualise the research. Chapter 3 describes the context and scope of this thesis. This will create a global overview of the actor landscape, research scope, and the theoretical framework within which this thesis is developed. Chapter 4 elaborates upon the methodology used for characterising the vulnerable devices and describes the interview protocols utilised for this research. Chapter 5 presents the results of the analysis of the vulnerable devices in the network of KPN. Moreover, chapter 6 entails the analysis of the end-users and manufacturer interviews. Chapter 7 discusses the results for a deeper understanding including the limitations of this study. Finally, chapter 8 concludes this thesis by answering the research questions and elaborating upon recommendations and future work.

2

Theoretical background

This chapter presents relevant previous scientific work. First, the security issues of internet-connected or IoT devices are discussed to provide background information on these devices (§2.1). Second, background on DDoS attacks and TCP reflective amplification attacks is presented (§2.2). Moreover, this chapter shows how vulnerable middleboxes and IoT devices can be used for such attacks. This background information provides insights into the potential impact and remediation methods of this vulnerability. Third, the current views in the academic literature on vulnerability notifications are shown (§2.3). Lastly, different views and models on people's behaviour in regard to cyber security are discussed (§2.4).

A Boolean search strategy is used to identify articles for this literature study. As certain combinations of keywords result in a large number of identified articles, the search strategy is divided into three steps:

1. Review of most recent literature
2. Review of most cited literature
3. Review of articles via snowballing (Jalali et al., 2012)

These three steps structure the approach of identifying relevant academic literature. By conducting step 1, the most relevant studies are analysed based on the number of citations. Step 2 will identify the current view and knowledge in regard to the topic. Lastly, step 3 adds detailed in-depth knowledge to the already found literature.

The comprehensive overview of the search strategy used to perform this literature study can be found in appendix A.

To improve the readability of this literature chapter and enhance the overall understanding of this thesis, a brief introduction is presented on internet communication. This introduction explains how internet-connected devices communicate with each other, which serves as the foundation for the main issue of interest in this thesis: TCP reflective amplification. This introduction is presented below:

Introduction to internet communication

The internet is a network of interconnected devices, such as computers and smartphones, that allow us to communicate with each other and access information from anywhere in the world. There are two main protocols used for internet communication: UDP and TCP.

TCP (Transmission Control Protocol) is a protocol that provides the delivery of data between applications running on different devices. Before any data can be transmitted between devices using TCP, a ”handshake” must take place between the sender and receiver. In a TCP handshake, the sender first sends an SYN (synchronize) packet to the receiver. The receiver responds with an SYN-ACK (synchronize-acknowledgment) packet. Finally, the sender sends an ACK (acknowledgment) packet to confirm that it has received the SYN-ACK packet. Once this three-way handshake is completed, the connection is established and data can be transmitted.

UDP (User Datagram Protocol), on the other hand, is a protocol that does not provide reliable delivery of data or ensures that data is delivered in the correct order. It does not require a handshake before sending data, which makes it faster and more efficient than TCP. However, UDP is susceptible to packet loss or duplication and it does not provide any mechanism to ensure that data is delivered reliably.

2.1. Security issues of internet-connected devices

Internet-connected devices or IoT have been a popular topic to study by academics. IoT presents various security issues which could be exploited by adversaries and cause harm to people. The main challenges for IoT security can be explained by two unique characteristics of IoT: the large scale of the devices and the heterogeneity of IoT devices (Panchiwala et al., 2020; Sicari et al., 2015; Z.-K. Zhang et al., 2014). First is the large scale of the devices, due to the large number of devices connected to the internet, IoT security suffers from scalability issues. It is practically impossible to protect and monitor every device all the time. Secondly, the term IoT is ambiguous. IoT devices can be everything with the only criterion of being connected to the internet. This makes IoT devices widely differ with regard to hardware and software specifications (Malina et al., 2016). Therefore, these devices are heterogeneous in nature. As a consequence, there is no ”one size fits all” security solution to protect all these devices from external threats.

These two characteristics can be discussed in more detail by looking at the following four security challenges with internet-connected devices. These four challenges describe the difficulties of protecting IoT devices when taking into account the two characteristics of IoT, in general, being heterogeneous and large in terms of the number of connected devices. These security challenges are based on a paper by Maple (2017). It must be noted here that these challenges are perceived as some of the most important challenges but these are not comprehensive for all IoT.

Physical limitations of devices and communication IoT devices have limited computation, memory and energy resources. That is why some security measures, such as encryption of data, are not feasible for all devices resulting in unprotected data. Additionally, due to these physical constraints, IoT often makes use of lightweight insecure communication protocols such as MQTT (F. Chen et al., 2020; Heer et al., 2011).

Authentication management Authentication within IoT devices is crucial as without proper authentication mechanisms in place, the confidentiality, integrity, and availability of the device can be compromised. Authentication in IoT is challenging because of the large number of different objects that comprise an IoT and therefore many internet-connected devices lack decent authentication mechanisms (Lin et al., 2017). Even when authentication systems are in place, devices can still easily get compromised because people do not make use of strong passwords or do not even change the default password. For example, the Mirai malware shows this as it was able to infect millions of IoT devices using default passwords (Yang et al., 2017).

Authorisation and access control Effective authorisation and access solutions found in conventional computer networks are hard to implement in the field of IoT. IoT access control systems need to be able to handle the scale of IoT, be easy to be managed due to the widespread of different applications of IoT devices and needs to be flexible to adapt to different users and their needs (Gusmeroli et al., 2013). Furthermore, conventional access control solutions such as role-based access and attribute-based access are known to be hard to implement in the IoT context because of low-powered IoT devices (Maple, 2017).

Updating, responsibility and accountability To protect internet-connected devices, it is important to address any found vulnerability as soon as possible before a device or system gets compromised. Simply updating or patching devices turns out to be difficult because some devices lack the capability to be updated or are too old for the manufacturer to produce patches for newly discovered vulnerabilities (Simpson et al., 2017). Even when security updates are available, for most security patches, users of internet-connected devices are needed to install the update or restart the device manually, which is challenging for IoT (De Carli et al., 2021). Furthermore, often it is unclear in the field of IoT who is responsible for securing the vulnerable devices and who is accountable in case an internet-connected device gets compromised (Maple, 2017).

2.2. DDoS attacks

Distributed Denial-of-service-attacks, or DDoS attacks, aim to disrupt the normal traffic of a server, service or network by flooding the target with internet traffic (Cloudflare, 2022). DDoS attacks are carried out by following various steps (Hoque et al., 2015; Kolias et al., 2017; Mirkovic et al., 2004). First, an attacker needs to recruit bots which can be used for launching the attack. A bot can be any internet-connected device. There are different ways to recruit these bots depending on the type of DDoS attack but it essentially means that the attacker exploits a vulnerability to gain (partially) control over an internet-connected device. After recruiting the device, the newly recruited bot can be remotely controlled by the attacker and used to launch an attack or recruit new bots to create a network of bots. This network of bots, or botnet, can then be controlled by the attacker to coordinate a continuous traffic stream to a potential victim so that legitimate users can not reach the services of the victim anymore. This process of carrying out a DDoS attack is visualised below in figure 2.

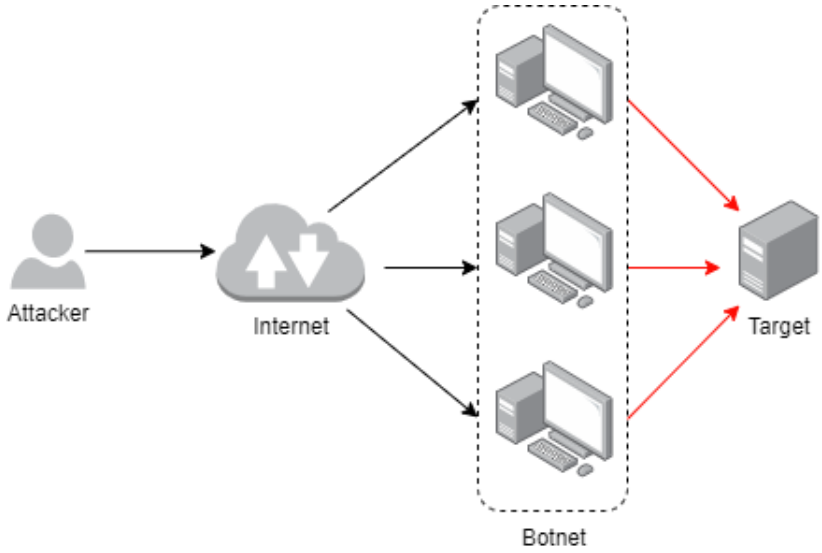

Figure 2: Process of DDoS attack

While all DDoS attacks work in general the same, as described above, there are multiple types of DDoS attacks which could be used by malicious people or organisations. DDoS attacks can be divided into three different types of DDoS attacks (Cloudflare, 2022; Das et al., 2015; Hoque et al., 2015). These types are based on the OSI-model, which is a conceptual framework which breaks down network communication into seven layers (Alani, 2014). Below, each of the three types is briefly discussed.

Application layer DDoS attacks Attacks on the application layer mainly target the HTTP protocol of a server with the goal to overload a web service. These types of attacks are hard to detect because it is difficult to distinguish legitimate from malicious traffic.

Protocol layer DDoS attacks Attacks executed over the protocol layer cause a service disruption by exploiting the weakness of communication protocols such as TCP. By flooding a server with TCP requests to a random port at the victims' server, a certain response is created by this device. In case of an SYN request flood, the attack exploits the TCP three-way handshake policy. By doing so, the victim server responds to all the SYN requests, exhausting its resources during the process.

Volumetric DDoS attacks Volumetric DDoS attacks try to disrupt the service of a victim by sending a large amount of traffic, consuming the bandwidth of the victim host. These large amounts of traffic are generated by some sort of amplification. This amplification occurs when the attacker sends a request to a vulnerable service running on a server. The service sends a huge reply package, causing server overload and service disruption. In case of a spoofed request[1], the large response will be directed towards the target IP address, disrupting the services of the victim. This type of attack is called an amplification reflection attack.

2.2.1. Reflective amplification attack

Amplification reflection attacks over services using the TCP protocol as the transport protocol are executed differently than attacks using services which are UDP-based. UDP protocols are interesting to be used for amplification reflection attacks because of the one-way traffic nature. The UDP protocol policy does not verify the IP address when responding to any request sent by an attacker, making it extremely easy for malicious people to abuse services that use this transportation protocol (Nuiaa et al., 2021). Examples of UDP-based services that are often abused for launching amplification reflection attacks are network service protocols such SSDP, SNMP, DNS, NTP or NetBios (Lyu et al., 2017; Rossow, 2014). The bots used for launching this type of DDoS attack do not have to be compromised, *reflection* can be created by sending a short query message to a vulnerable service with a spoofed IP address, to which the vulnerable devices responds by sending a large response to the victim (Lyu et al., 2017). This process is illustrated on the left side in figure 3.

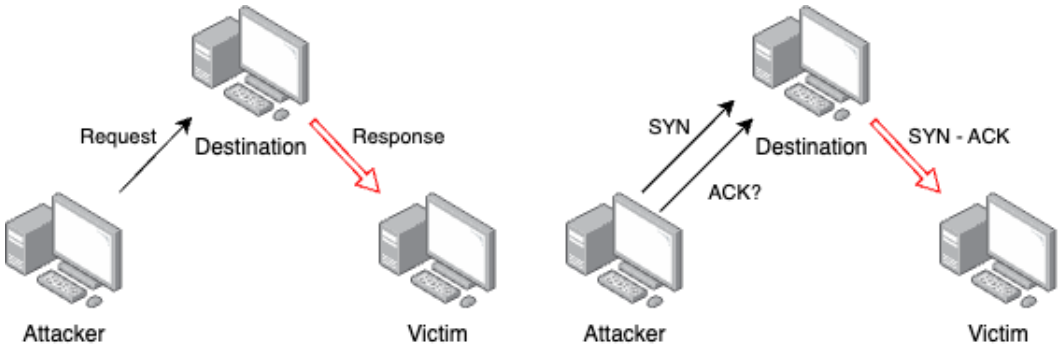

Figure 3: Amplification reflection attack over UDP-based services (L) and TCP-based services (R)

The three-way handshake principle prevents attackers to use TCP-based services to launch effective amplification attacks. Although sending packages to TCP-based services result in some reflection, it does not allow for easy amplification as it does for UDP-based services (Rossow, 2014). The three-way handshake allows the receiver of the traffic to verify the IP address of the sender.

Nonetheless, this protocol can still be used for amplification reflection attacks. TCP amplification reflection attacks work according to the following two steps (Bock et al., 2021; Kührer et al., 2014b). First, the attacker sends an SYN request to the destination device with a spoofed IP address of the victim. The destination device responds by sending the SYN-ACK packet to the victim, resulting in some amplification. Because the attacker does not have received the SYN-ACK packet, the three-way handshake can not be completed and no full connection can be made. Because the handshake is not completed, the attacker is not able to make a connection to send a request to the destination needed for creating large

[1] Falsifying the source IP address in a packet to disguise the identity of the sender or to launch an attack.

amplification as is possible on UDP-based services. Therefore, only a small amplification is created by the SYN-ACK packet. This process is demonstrated on the right side in figure 3.

2.2.2. Abusing middleboxes or IoT for TCP reflective amplification attacks

The paper by Bock et al. (2021) is the first to explain how middleboxes can be abused for launching large-scale cyber attacks. As described in this paper, these vulnerable middleboxes can be used to launch DDoS attacks.

The key understanding, in how to use vulnerable middleboxes for launching TCP reflective amplification attacks, is to not aim for a response from the destination devices, but rather from a middlebox which is on the path to the destination (Bock et al., 2021). Figure 4 shows the process of launching such attacks on the left side. First, the attacker sends a packet to a certain server with the spoofed IP of the intended victim. Often these packets go through a path of various devices, including middleboxes. A middlebox can be any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host (Carpenter et al., 2002). Examples of middleboxes are firewalls or intrusion detection systems (IDS). According to Bock et al. (2021) these middleboxes have two characteristics which make them interesting for this attack. First, middleboxes expect to miss some packets and therefore can assume that the three-way handshake has been completed. In case they see a request for a forbidden URL (such as websites for gambling or pornography) a middlebox can respond without a valid TCP connection. The second characteristic is the fact that middleboxes often inject traffic into the destination host to block the connection when a forbidden request is sent. This response by the middleboxes allows for amplification as it sends a large block page to the victim, as illustrated in figure 4.

TCP reflective amplification attacks via vulnerable IoT devices operate in a similar manner to those that exploit vulnerable middleboxes, but with some notable distinctions. The principal difference is that amplification is generated by the destination device instead of a device en route to the destination. This process is visualised on the right side in figure 4 and looks as follows. The attacker transmits a packet to an IoT device with a spoofed IP address of the target victim. This packet contains a request for an URL, which can be a conventional URL or a restricted one similar to middleboxes. In normal circumstances, a device would not respond to the request as the TCP handshake cannot be completed. However, if the TCP protocol is configured improperly on the device, it can respond to the request with the HTML landing page of the IoT device, disregarding the fact that the handshake is incomplete. Unlike middleboxes, it does not respond to the restricted URL request but replies to any URL request by providing the landing page of that particular device, generating substantial amplification.

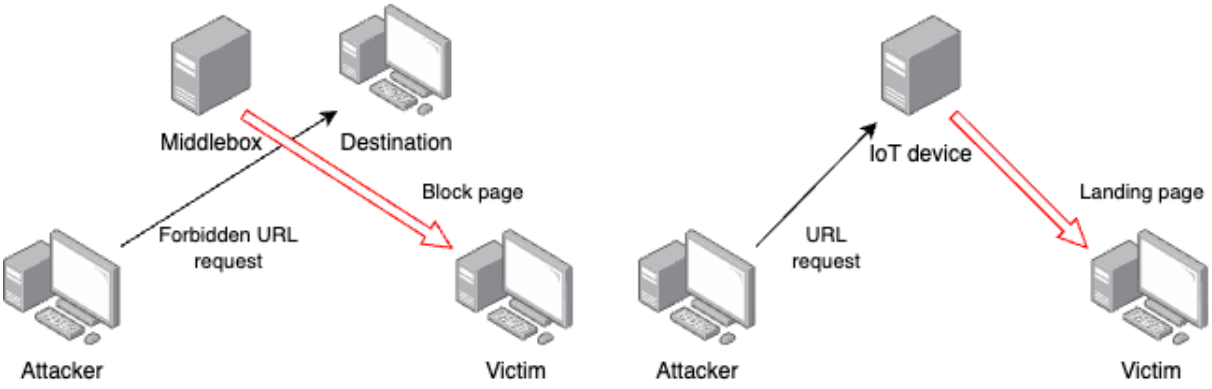

Figure 4: TCP reflective amplification attack via middleboxes (L) and IoT devices (R)

Bock et al. (2021) identify in their paper millions of IP addresses which can act as amplifiers for launching TCP amplification attacks using the principle as described above. They concluded that 82.9 per cent of these identified IP amplifiers are actual middleboxes by looking at the traceroutes to the IP addresses, suggesting that the remainder are other types of internet-connected devices. Furthermore, they found exceptionally high amplification factors, which often exceed the well-known UDP-based amplifiers. The main causes for these large amplification factors are victim sustains loops, where the victim host sends a reset packet back to the destination as this device did not expect the block page. Some middleboxes

respond to this reset packet with another block page, creating a loop which results in lots of traffic generated towards the victim. The other main cause identified are routing loops. When a vulnerable middlebox is positioned in this loop, every time a forbidden request passes this middlebox, a response is generated with a block page that is sent to the victim.

2.2.3. Remediation of vulnerable middleboxes

The issue of vulnerable middleboxes can be solved but it is not easy as there is not a single vendor or network that can be patched to correct the problem (Bock et al., 2021). The same can is likely the case for IoT devices, which can be any device with a broken TCP protocol. The four solutions provided by Bock et al. (2021) focus on changing the configuration of middleboxes:

- Require middleboxes to see traffic in both directions instead of one in order to check if the TCP handshake has been completed. This solution aims to make it more difficult for attackers to spoof connections by requiring middleboxes to see traffic in both directions before sending large block pages to the victim.
- Limit the size of injected responses by middleboxes. To prevent middleboxes from enabling amplification attacks, this recommendation suggests limiting the size of injected responses or using alternative methods such as RST packets or minimal HTTP redirects instead of large block pages.
- Configure middleboxes to only censor requests originating from within the intended network. This solution is focused on the many state censorship middleboxes. This can be solved by limiting the scope of victims of amplification attacks by configuring middleboxes to only censor requests from within the intended network or country.
- Remove or limit censorship devices. This solution suggests removing or disabling HTTP injection in outdated censorship devices to prevent abuse by attackers and reduce the potential for amplification attacks.

Additionally, both Pal (2022) and CIRT Akamai (2022), note that SYN packages usually contain no data thus SYN floods with package sizes greater than 0 can be considered suspicious. Vendors and manufacturers are often responsible for solving the issue of these middleboxes and IoT devices as they can alter the configuration by releasing security patches. At the moment of writing, there are already a few examples of middlebox vendors releasing security patches for this issue (Pal, 2022). Therefore, it is important that end-user of vulnerable middleboxes and IoT devices regularly check vendor advisories for new security patches.

2.3. Vulnerability notifications

Vulnerability notifications are a suggested technique to inform end-users of vulnerable or malware-infected devices to remediate the vulnerability. For example, F. Li, Bailey, et al. (2016), conducted a study by notifying thousands of different network operators of security issues in their network. The study included three different vulnerabilities including accessible industry control systems, misconfigured firewalls and hosts that can act as DDoS amplifiers. They concluded that vulnerability notifications have a positive effect on patching when network operators are contacted directly with detailed information about the vulnerability. Stock, Pellegrino, Rossow, et al. (2016) concluded the same from their research but showed that they encountered difficulties in reaching operators of vulnerable hosts. J. Zhang et al. (2017) tries to address this issue by researching the topic of vulnerability notifications from an ISP perspective, as ISP's have the contact information of their customers which makes it possibly easier for them to reach end-users than for instance researchers. This study identified instant messenger to be the most effective in terms of vulnerability clean-up rate compared to the two other tested methods of notifying customers via mail or phone call. Nevertheless, they state that instant messenger is not commonly used by IPS's and that the effectiveness is influenced by not only the communication channel but also the notification content, the vulnerability characteristics, the technical capability of the operator and the level of authority of the operator to perform the vulnerability cleanup. Another study by Durumeric et al. (2014) looked into the effects of vulnerability notifications send to network operators of websites which were vulnerable to the Heartbleed vulnerability. The results of this research indicated that vulnerability notifications have a positive influence on the patching rate of vulnerable devices. The patching rate of Heartbleed-infected devices increased by 47 per cent after sending a direct

vulnerability notification to the network operator (Durumeric et al., 2014).

Besides the importance of the notification channel and the notification content, another factor identified by researchers influences the effectiveness of vulnerability notifications. Multiple studies revealed that there is a high level of distrust by receivers of vulnerability notifications and that the reputation of the senders seems to play an important role in the remediation rates (Çetin et al., 2016; Hennig et al., 2022; Rodriguez et al., 2022; Stock, Pellegrino, Backes, et al., 2018; Zeng et al., 2019). Stock, Pellegrino, Backes, et al. (2018) show that receivers of vulnerability notifications perceive low levels of trust in either the notification itself or the sender of the notification. The same observation is later also discussed in the study by Hennig et al. (2022). Interestingly, results from a study by Çetin et al. (2016) show no evidence that the reputation of the sender of vulnerability notifications influences the remediation rate of a vulnerable device. van Eeten (2017) draws the same conclusion from his research, but it must be noted here that for this study vulnerability notifications were sent to ISP's instead of to consumers. A different notification perception by ISP's and consumers may explain why different studies draw different conclusions.

Overall, academics agree that the communication channel and notification content both play an important role in creating effective vulnerability notifications. Therefore, the following two sub-chapter explore the topic of notification channels and content in-depth.

2.3.1. Notification channels

Previous research studied various possible notification channels which could be used by ISP's to send their vulnerability notification. It is up to the ISP which notification mechanism to adopt. Livingood et al. (2012) have published an article with recommendations on how an ISP can use various communication channels to contact their customers in order to increase the remediation rate of vulnerable or malware-infected devices. While this article is focused on notification channels that can be used for vulnerability notifications to remediate malware-infected IoT, these channels can be applied to notify end-users of devices with other vulnerabilities as shown by the discussed research articles in chapter 2.3. Below, each of the notification mechanisms, named by Livingood et al. (2012) is briefly discussed. Some of these channels are discussed more in-depth than other because of their relevance to this research.

Email notification Sending vulnerability notification over email, is the common way to contact customers of an ISP (Livingood et al., 2012). This is an easy-to-adopt and cheap method for warning ISP customers about possible vulnerabilities. A possible drawback is having no guarantee that the receiver reads the mail or that the receiver perceives the mail as spam or phishing. Notifying end-users via email, and the impact of such email is researched before by various studies such as Bouwmeester et al. (2021) and Cetin, Ganan, Altena, Kasama, et al. (2019). While this type of communication channel is covered by multiple academics, there is no consensus on the effectiveness in terms of vulnerability clean-up rate. For example, Cetin, Ganan, Altena, Kasama, et al. (2019) found no evidence that mail-only notification has an impact on mitigation measures compared to a control group which did not receive a notification. However, F. Li, Bailey, et al. (2016) did conclude that email notifications had a positive effect on the cleanup efforts of the vulnerability. A possible explanation for these contradictory findings could be due to the different vulnerabilities researched in the different studies.

Telephone call notification Livingood et al. (2012) state that telephone call notifications may be an effective way, in particularly high-risk situations. Nevertheless, they explain that for ISP's phone calls are often not feasible due to the high costs of making a large number of calls in terms of financial and organisational resources. This observation is validated by Altena (2018) during her research at a Dutch ISP.

Postal mail notification The least common way of notifying as it is assumed to be ineffective due to delivery time and cost.

Walled garden notification Walled garden notification places a vulnerable host in a restricted environment that controls the information and services that the host is allowed to make use of (Livingood et al., 2012). It is found that walled garden notifications are one of the most effective in terms of malware cleanup rate compared to other notification channels (Cetin, Ganan, Altena, Tajalizadehkhoob, et al., 2019)

Instant message (IM) notification IM message notifications provide ISP with a simple way of notifying their customers. This method is both cost-effective and easily scalable (Livingood et al., 2012). Furthermore, J. Zhang et al. (2017) show that instant message notifications resulted in the highest vulnerability clean-up rate among the tested types of notification channels. However, there are some disadvantages that can be named to this method. For example, not every ISP customer uses IM and there might be privacy concerns when utilising a third-party IM service for sending notifications.

Short message service (SMS) notification Similar to IM notifications, SMS notifications are an easy way of notifying ISP customers. A major disadvantage is the limited amount of characters which can normally be used in such notifications. Therefore, SMS notifications could be ineffective in some cases because not all necessary information can be conveyed which is needed to perform the vulnerability cleanup.

Web browser notification Web browser notifications are intended to notify devices running a web browser, making it not applicable for most IoT devices.

All of the above-named communication channels have their advantages and disadvantages, therefore it is suggested to combine different communication channels in order to achieve high levels of vulnerability clean-up rates (Livingood et al., 2012).

2.3.2. Notification content

The content of a vulnerability notification plays an important role when trying to reach high levels of vulnerability remediation rates. The content should be clear, easily understandable and incentivising to perform certain actions. Previous studies state that vulnerability notifications should be comprehensive in order to be effective (Çetin et al., 2016; Durumeric et al., 2014; F. Li, Ho, et al., 2016). While other researchers state that notifications should be simple and easy to interpret (Forget et al., 2016).

The content should be tailor-made for specific end-user in order to be effective. Rodriguez et al. (2022) explain that it is important to understand the technical abilities of the end-users who receive vulnerability notifications in order for them to comprehend the notification, understand what actions need to be performed and know how to do it in a sufficient way. Rodríguez et al. (2021) conclude from the results of their study that, once end-users get notified of a security problem, they are willing and able to remediate the vulnerability, irrespective of the comprehension of the information in the notification. As an explanation for this behaviour, they illustrate that this effect could be due to early IoT adopters being also more technically competent than users in general. As IoT becomes more widespread, the notification content could have a bigger impact on the remediation rate as IoT users are less tech-savvy than before (Rodríguez et al., 2021).

2.4. End-user security behaviour

To understand why end-users perform certain actions after receiving a vulnerability notification, it is important to understand their motives and perception. In the past, various literature discussed the topic of consumer security behaviour and their actions after receiving a vulnerability notification or a security warning. The motives and perceptions of the end-users in regard to security behaviour can be discussed based on different models, including the C-HIP model, COM-B model and PMT (Agrawal et al., 2020; Menard et al., 2017; Rodriguez et al., 2022). These models can be used for understanding why and when people do comply (or not) with recommended security measures communicated via vulnerability notifications. Below each of these three models is discussed in more depth in combination with some examples of using these models for research in vulnerability notifications.

PMT Protection motivation theory (PMT) was originally created by Rogers (1975). The theory explains how people respond to threats. It states that when an individual is confronted with a threat, they will go through a process to assess the severity of the threat and the availability of coping resources. Based on this assessment, they will choose to either take action to protect themselves (adaptive response) or avoid taking action (maladaptive response) (Rogers, 1975). Adapted versions of this theory have been applied in the domain of cyber security in order to assess what motivates people to comply with security policy or to explain people's tendency to engage in voluntary secure behaviours (Boehmer et al., 2015; Menard et al., 2017; Tsai et al., 2016).

According to PMT, behaviours are motivated by the evaluation of both threats and coping factors (Rogers, 1983). Threat appraisals are determined by one's perception of the vulnerability to risk and the perceived consequences of unsafe behaviours. Coping appraisal involves assessing one's belief in their ability to successfully execute protective behaviours (coping self-efficacy), the perceived effectiveness of those behaviours (response efficacy), and the response costs of enacting those behaviours. Together, these two appraisals influence an individual's intentions to adopt protective measures (security intentions) and show certain behaviour. This process is shown in figure 5.

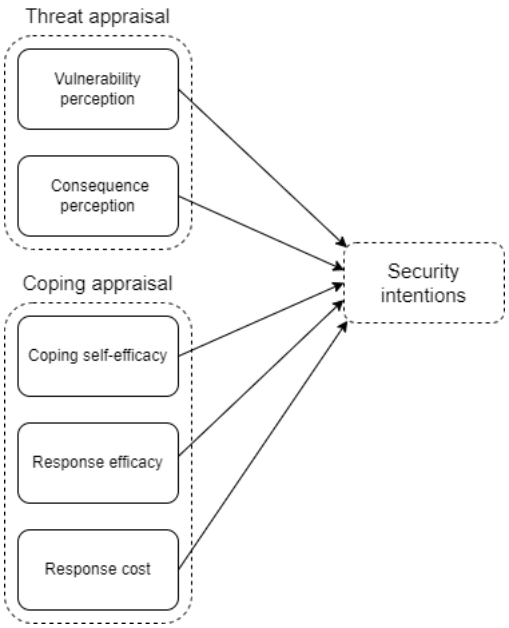

Figure 5: PMT model based on Rogers (1983)

C-HIP model Academics have used the Communication-Human Information Processing (C-HIP) model for describing end-user behaviour when receiving vulnerability notifications or security warnings (Agrawal et al., 2020; J. Chen, 2020; Rodríguez et al., 2021). Developed by, Conzola et al. (2001), this model structures the stages of security warning by describing seven stages between the sender of a warning and the receiver of the warning who is expected to process the information and produce behaviour. These stages can be used for organising the findings in warning research and the model can serve as a tool for determining why a security warning is ineffective in its goal of influencing the behaviour of an end-user (Wogalter, 2006). The seven stages, according to Wogalter (2006), can be described as follows:

- **Delivery:** describes whether the notification or warning is received by the intended recipient
- **Attention switch:** the process where the recipient moves his attention towards the notification from some other activity
- **Attention maintenance:** attention must be maintained after switching the attention to the notification in order to acquire all the information in this notification
- **Memory:** after acquiring the information, the recipient must try to comprehend the new information and store it in a memory
- **Beliefs:** the content of the notification must be in line with what the receiver believes is true such as the risk perception of the mentioned threat in the notification
- **Motivation:** the receiver needs some form of motivation in order to comply with the measures in the notification
- **Behaviour:** the ultimate measure of the effectiveness of the notification

Other important aspects to consider, which are mentioned in the C-HIP model, are the source of the notification and the channel used for communicating the notification. Furthermore, there are other environmental stimuli that can influence how the warning is processed by the receiver such as other people or other warnings. The feedback loops show that later stages can influence processing in earlier stages as Wogalter (2006) explains that the model is not linear. When the information from the notifications flows through the whole system, and every stage is passed, the receiver will change its behaviour. A complete overview, including the seven stages, can be found below in figure 6.

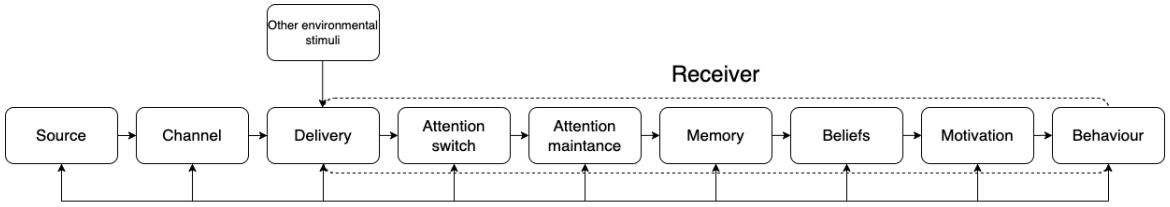

Figure 6: C-HIP model based on Wogalter (2006)

Fagan et al. (2015) use the model for analysing the results from their study into the end-user perception of software update notifications. It was found that end-users are often annoyed and sometimes confused by update messages, resulting in the undesired behaviour of not updating the software. Research by Rodríguez et al. (2021) studied the effect of vulnerability notifications on malware cleanup and performed a phone survey according to the C-HIP model. The results from this study suggest that almost every participant took some action after receiving a vulnerability notification thus showing behaviour. Nevertheless, only 24 per cent of all survey respondents were able to perform all malware remediation steps as provided in the notification (Rodríguez et al., 2021).

Other research by Ormond et al. (2022) proposes in their study a combination of the C-HIP model with the PMT theory to capture the effect of security notifications on the behaviour of end-users. Based on this literature study they conclude that at the moment the perception, motives, attitudes and beliefs of security warnings are understudied while the focus is most of the time on notification channels and compliance behaviour.

COM-B model Rodriguez et al. (2022) interviewed end-users of malware-infected devices using a design informed by the COM-B model. The Capability, Opportunity, Motivation and Behaviour (COM-B) model, designed by Michie et al. (2011), explains how the three essential conditions capability, opportunity and motivation interact to generate a certain behaviour which in turn affects the three components. This model is a useful tool for both understanding, as well as predicting, behaviour (Michie et al., 2011). Capability is defined as the capacity of an individual to engage in the activity as presented in the notification. Motivation refers to the psychological factors that drive behaviour, such as attitudes, beliefs, and values. Opportunity describes the external presence of the necessary resources and information needed to perform a behaviour.

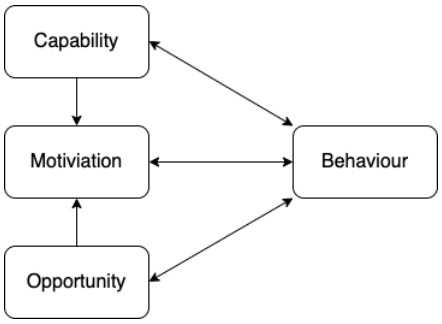

Figure 7: COM-B model based on Michie et al. (2011)

Rodriguez et al. (2022) explain that these three conditions act as pillars which all need to be in place to create the context for a behaviour change intervention to be considered complete. Furthermore, this research states that the COM-B model is suitable for analysing the behaviour of consumers after

they receive a vulnerability notification. This behavioural model is also used in a study by van der Kleij et al. (2021), which focused on cyber security behaviour in a more general view. This study concluded, based on a survey conducted in The Netherlands, that people with relevant knowledge, and who are motivated to protect themselves against cyber threats, report more cyber secure behaviour. The opportunity people have to protect themselves was found to be only partly related to more secure behaviour.

2.4.1. Patching behaviour

Patching is a proposed solution to cope with the issue of vulnerable middleboxes, as explained in chapter 2.2.3. The vulnerability life cycle, as explained by Arbaugh et al. (2000) and Frei et al. (2006), can be used to explain the patching behaviour in regard to vulnerabilities. This life cycle consists of three key moments in time: discovery time, patch time and disclosure time. Figure 8 shows how the number of attacks resulting from a software vulnerability changes over time. It starts when the vulnerability is created during the software development stage. Then, when the vulnerability is found, some small-scale attacks by exploiting the vulnerability may occur. However, most exploits happen after the vulnerability is publicly disclosed. The number of attacks gradually decreases after a patch is released, and the vulnerability's life cycle ends when all affected systems are patched.

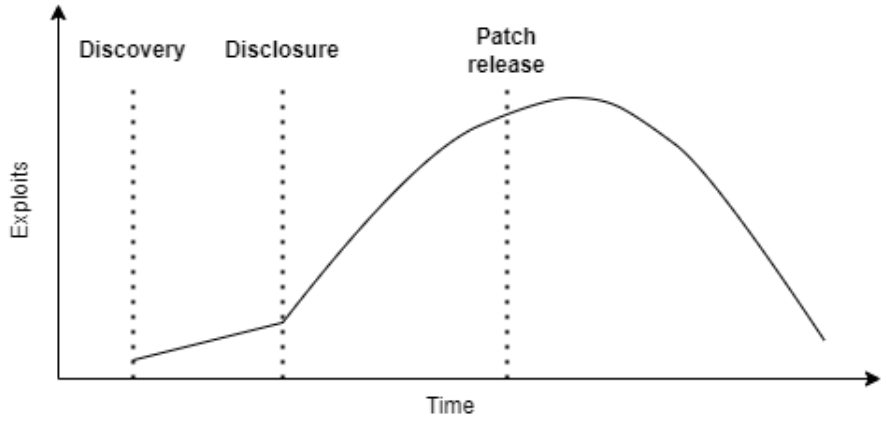

Figure 8: Vulnerability lifecycle based on Arbaugh et al. (2000)

Discovery The time of discovery is when a software vulnerability is first identified as a security risk. A software vulnerability exists from the moment the software is developed, but it remains dormant until it is discovered. It may not be made public until after it is disclosed, and in some cases, the discovery date may never be publicly known.

Exploitation A vulnerability becomes problematic once an exploit is available, and exploits increase after the disclosure of a vulnerability. The time of exploit is the earliest date that an exploit for a vulnerability becomes available. An exploit is described as any tool, virus, data, or sequence of commands that takes advantage of the vulnerability.

Disclosure The time of disclosure is the first date that a vulnerability is publicly discussed on a channel by a trusted and independent source that undergoes expert analysis and includes risk rating information. Making a vulnerability public increases attackers' knowledge of it, which is why most exploits occur after disclosure.

Patching Releasing a patch marks the start of the decay of a vulnerability, but the patching process can take a long time, especially for larger vendors that require planning and testing. Moreover, existing literature indicates that the patch release behaviour is influenced by the vulnerability disclosure, stressing the importance of vulnerability disclosure (Sen et al., 2020).

End-user security patching behaviour has been studied in the past. For example, Haney et al. (2021) researched smart home user perception of updates. They found that users were often confused about how and if security patches have been applied. Furthermore, they state that users rarely link updates with security. Other research, looking at the lifespan of exploits in IoT malware, identified the lack

of patching as one of the main causes for the vulnerabilities (Al Alsadi et al., 2022). This research suggests educating users on the importance of patching IoT devices. In general, scholars agree that users frequently fail to update their devices within an appropriate timeframe or neglect to update their susceptible devices altogether due to various factors, such as limited awareness of associated risks and the trade-offs between not being able to use the device during the update process. (Rajivan et al., 2020; Sarabi et al., 2017).

3

Research context

This chapter describes the environment where this research is conducted. First, a stakeholder analysis is performed to sketch the actor field in which the problem of TCP reflective amplification attack is situated (§3.1). The second sub-chapter discusses the research scope, which includes providing background information on the ISP, which facilitated the tools and information needed for carrying out the analysis for this thesis (§3.2). Last, a theoretical framework is presented which is used for analysing the vulnerable device end-user behaviour (§3.3). Together, this gives an insight into the context where this research has taken place and what aspects have been included or excluded from this research.

3.1. Stakeholder analysis

The security of IoT devices consists of a wide range of inter-connected actors (Kar et al., 2018). Therefore, it is important to understand which actors are involved, and their role, when exploring the issue of vulnerable middleboxes and IoT devices which can be abused for TCP reflective amplification attacks. Below, in figure 9 the different actors involved are displayed together with their inter-connected relations. Not all involved actors are shown. For example, governmental institutions are involved as they are responsible for creating policies, regulations and laws which apply to IPSs and manufacturers. However, because this is out of the scope of this research, this actor is not included in the figure. For more information on the scope of this thesis, see chapter 3.2.

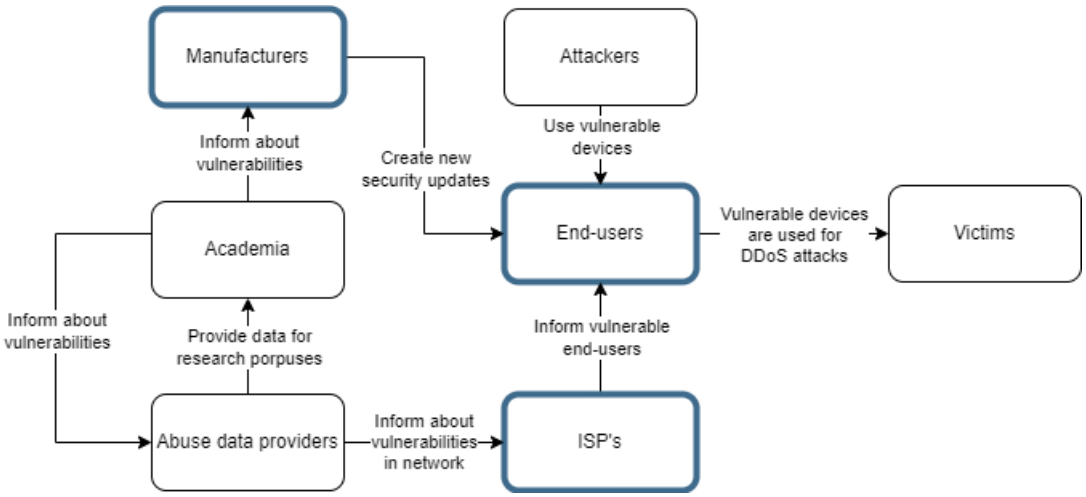

Figure 9: Overview of the actor field

This research focuses on three key actors, namely KPN (ISP), vulnerable device end-users and manufacturers, which will be studied in-depth. The following paragraphs provide a detailed elaboration of their respective roles.

ISP's ISP's are accountable for providing and maintaining the internet infrastructure, which enables both end-users and businesses to access the internet. They bear the responsibility of addressing security concerns and vulnerabilities in their network that could adversely affect their customers or the performance of their network. Third-party organisations such as Shadowserver provide ISP's with information regarding such security issues and vulnerabilities. ISP's occupy a unique position in that they can detect vulnerabilities within their network and possess the contact information of their customers, allowing them to inform them of any detected vulnerabilities.

End-users End-users own the vulnerable devices and, despite being unaware of the device's vulnerability or abuse, are considered victims. They have the ability to take action to address vulnerabilities, for example by installing updates or disconnecting the device from the internet. However, as they bear the burden of such responsibility, they are also susceptible to negative consequences making them possible victims of IoT abuse.

Manufacturers Finally, manufacturers or vendors who design and produce internet-connected devices form the third key actor in this study. These actors are responsible for designing and implementing security updates in the event of vulnerabilities in their devices. Furthermore, they have a communicative role towards their customers, notifying them of any new software or hardware updates.

3.2. Research scope

This research analyses the end-users, responsible for vulnerable middleboxes and IoT, and their devices in the network of KPN. KPN is one of the largest ISP's in the Netherlands with a market share of 35 to 40 per cent, with only VodafoneZiggo being larger with a market share of 40 to 45 per cent (ACM, 2023). Therefore, it might be possible to generalise some of the findings for the whole population in the Netherlands.

Furthermore, the focus of this research is on the consumer market (CM) of KPN. The business market (BM) and wholesale market (ZM) are excluded from this research, as these are notified about abuse incidents on a best-effort basis and are not included in the same notifying process as customers in CM. Besides, businesses possibly need a different approach to solving the issue of vulnerable devices and behave differently in terms of reactions and actions after receiving a vulnerability notification. Making them less applicable for including them in an experiment with consumer end-users. This is further discussed in chapter 3.2.1.

All consumer IP addresses with devices vulnerable to TCP reflective amplification attacks are included in this research. This concerns IP addresses reported by an external party between the 25th of April 2022 and the 19th of February 2023. A period of 5 weeks, between the 16th of January and 19th of February 2023, is considered to be the research period. Every IP address reported during this time frame is considered vulnerable during the research period, the other IP addresses are considered to be not vulnerable anymore during this research. This differentiation is explained in depth in chapter 4.

3.2.1. KPN abuse department

Within, KPN, the abuse department is responsible for handling abuse incidents on the KPN network. The abuse department is a subsidiary of the larger KPN blue team, which is part of the KPN Chief Information Security Office (CISO). The primary goal of the abuse department is to remediate vulnerabilities which are present at their customer's IP addresses and handle incidents where customer IP addresses are abused for various kinds of malicious activities. Abuse could be intentional or unintentional, which is why the abuse team is also tasked with notifying customers about vulnerabilities and possible abuse of customer IP addresses.

The abuse department receives data from third-party organisations and other external sources on security incidents, vulnerabilities and other malicious activities within the KPN autonomous system (AS). The abuse feeds can be divided into two distinct sources;

- Shadowserver and other nonprofit organisations send daily abuse feed via email with data to the KPN systems which process this data.
- Individuals that report abuse incidents. This can only be done via email, the abuse department

can not be contacted via telephone or other means of communication due to the limited amount of available resources.

The data in regard to vulnerable middleboxes and IoT that can be used for TCP reflective amplification attacks, which is applicable to this research, is reported on by Shadowserver. This data is discussed in depth in chapter 4.

The abuse feeds, as described above, are automatically processed and vulnerability notifications are sent over email to the customer linked to the IP address. This process is completely automated. Customers can reply to the email for further assistance from the abuse department for remediating the vulnerability, malware or virus detected. In case of severe security incidents, such as malware infection, KPN uses so-called ”walled garden” notifications where customers' IP is placed in a quarantine environment where they can solve the issue detected. Customers can be released from the quarantine environment by filling out the contact form, by the manual release by the abuse department or after the expiration date of the quarantine passes.

Prior to this research, KPN did not process the data from Shadowserver or any other source, on middleboxes or IoT devices vulnerable to TCP reflective amplification. Therefore, no quantitative or qualitative data is available from prior experiences on notifying KPN customers about vulnerable middleboxes.

3.3. Theoretical framework

For this research, the COM-B model formulated by Michie et al. (2011), is used to analyse the behaviour of end-users who own a vulnerable middlebox or IoT device which can be abused for TCP reflective amplification attacks. Even though the model is not widely used for describing end-user behaviour in cyber security, the COM-B model has been proposed as a tool to understand the motivators and barriers to end-user security behaviour (ENISA, 2019). In chapter 2.4, this model is already briefly discussed. This sub-chapter will elaborate more on the model and will explain how it can be applied to vulnerability research in this thesis.

The COM-B model consists of three different components which interact with each other (West et al., 2020). These are capability, opportunity and motivation which describe the last component of the model: behaviour.

Capability Michie et al. (2011) describes the first component, capability, as the capacity of an individual in terms of psychological and physical capacity that makes it able for an individual to perform a certain behaviour. It describes that an individual needs the necessary skills and knowledge to perform a certain behaviour. Psychological capacity is for example the individual's cognitive ability and emotional state, while physical capacity refers to the individual's physical ability and health status. Capability also includes the individual's self-efficacy, which is the belief in one's ability to perform a behaviour successfully.

Opportunity The second component in the COM-B model is opportunity, which can be divided into physical and social opportunities that enable behaviour. Physical opportunities refer to the physical environment in which a behaviour occurs, including access to resources, facilities, and equipment. Social opportunities, on the other hand, refer to the social context in which behaviour occurs, including social norms, expectations, and support from others. In addition, the opportunity component also considers the availability of time and financial resources, which may impact an individual's ability to engage in a behaviour.

Motivation The third component in the COM-B model is motivation, which refers to the brain processes that direct an individual towards a particular behaviour. This component can be divided into reflective processes and automatic processes. Reflective processes involve cognitive and deliberative thinking, such as beliefs, attitudes, and intentions that guide behaviour. Automatic processes, on the other hand, are instinctual and impulsive, such as emotions, habits, and conditioned responses that influence behaviour. Additionally, motivation is influenced by an individual's goals and the level of willingness to perform and show a certain behaviour.

In turn, behaviour also influences the three components of the COM-B model. Engaging in behaviour

can improve an individual's capability by enhancing their knowledge and skills. It can also create new opportunities for behaviour by changing the physical and social context in which behaviour occurs. Moreover, behaviour can reinforce or weaken an individual's motivation by creating positive or negative experiences that influence their beliefs, attitudes, and emotions towards the behaviour.

Therefore, the COM-B model proposes that behaviour is influenced by the interplay between capability, opportunity, and motivation and that behaviour also has a mutual influence on these components. Understanding the complex interrelationships between these components can help design effective behaviour change interventions that address the root causes of behaviour.

This model will be used for analysing the behaviour of end-users with vulnerable middleboxes or IoT devices which can be abused by TCP reflective amplification attacks. Chapter 4.4.1 demonstrates how the COM-B model is implemented in the end-user interview protocols. Furthermore, chapter 6.3 links the results of these interviews back to the constructs of this model.

4

Methodology

This chapter presents the methodology employed to answer the main research questions and achieve the goals of this thesis. Firstly, the different data sources utilised for this research are discussed in detail (§4.1). Secondly, the quantitative data used in this thesis is discussed (§4.2). Third, the end-user interview setup is elaborated upon (§4.3). The subsequent subsections focus on the interview protocols employed for conducting end-user (§4.4) and manufacturer interviews (§4.5). The chapter concludes by discussing the limitations of the method (§4.6) and ethical considerations (§4.7) involved in this research.

4.1. Data sources

In order to answer the main research question and the sub-questions, different data sources will be called upon and various methods will be used for obtaining and analysing the data. Therefore, it is important to provide a clear overview of how these data sources will be used in this thesis as this will give structure to the methodology of this research in general. Figure 10 presents an overview that illustrates the various data sources and their interactions.

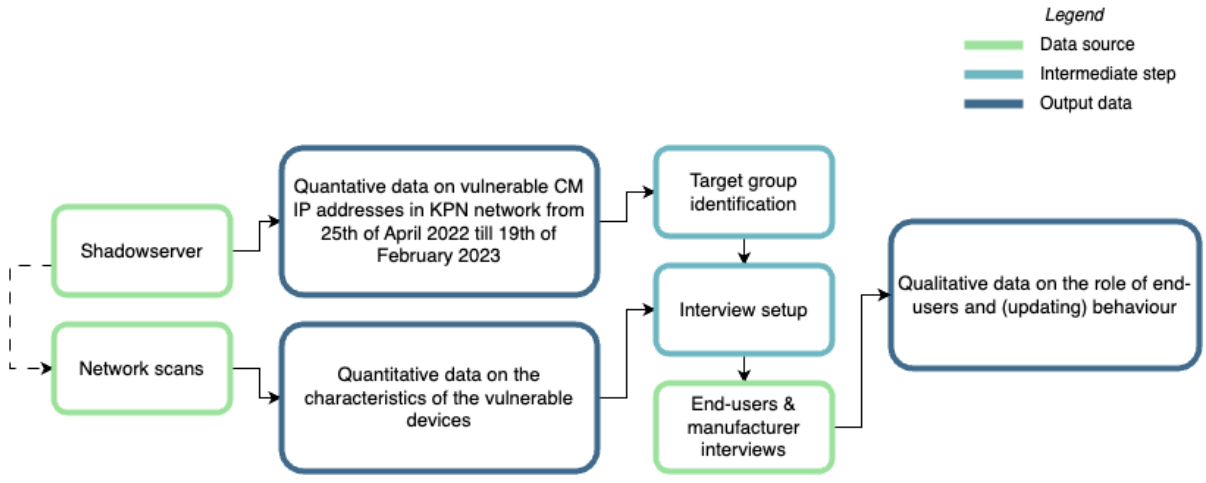

Figure 10: Overview of the data used for the research

This research utilises data from three distinct sources. First, quantitative data on vulnerable middle-boxes or IoT devices is provided by a third party called The Shaddowserver Foundation. The second source is quantitative data obtained through network scanning of IP addresses with vulnerable devices, as reported by Shadowserver. The third source is qualitative data directly provided by vulnerable device end-users and a manufacturer through interviews.

The Shadowserver data serves as the starting point for this research. It provides insights on vulnerable IP addresses in the KPN network from April 25th 2022 to February 19th 2023. This data is used to identify three distinct target groups. Additionally, the IP addresses reported by Shadowserver are used as input for the IP scans, which make up the second data source in this thesis.

The network scans provide quantitative data on the characteristics of the vulnerable devices, which, combined with the identified target groups, were used to set up an experiment. In this experiment, vulnerable device end-users were notified of the vulnerability via a vulnerability notification. Based on this setup, end-users and a manufacturer are interviewed to obtain further qualitative data on the role of the end-users and the manufacturer.

All three types of obtained data are analysed and used to answer the main research questions and sub-questions. The following sub-sections will discuss the three data sources, as illustrated in figure 10.

4.1.1. The Shadowserver Foundation

The Shadowserver Foundation reports on a wide variety of vulnerabilities, malicious activities and emerging threats on the internet with as main goal to make the internet more secure (Shadowserver, 2023). The services provided by this organisation assist companies globally, including ISPs, in identifying and responding to vulnerabilities within their network or organisation. Additionally, they support researchers in studying cybersecurity threats.

Shadowserver started daily reporting on vulnerable middleboxes in April 2022 (Shadowserver, 2022). The scans are performed in a similar way as described in the paper by Bock et al. (2021). They identify abusable middleboxes through scanning with custom TCP packets. The scans are executed by sending the following two different custom TCP packets to port 80:

- SYN packet with an HTTP GET payload for a forbidden URL (such as pornography or gambling sites)
- SYN packet with sequence number x, followed by a PSH + ACK packet with sequence number x+1 containing an HTTP GET payload request for a forbidden URL

The daily reports provided by Shadowserver allow this research to track IP addresses with vulnerable devices over time. These reports include the observed amplification factor, IP address and some basic geographical information and are filtered so that it only contains IP addresses of KPN customers. The reports can give an indication of whether the vulnerability is mitigated as the IP address will stop being reported on in the logs. It is important to note that while Shadowserver scans for vulnerable middleboxes, the results of this research suggest that the reported IP addresses by Shadowserver do not exclusively pinpoint vulnerable middleboxes, but also include IoT devices with a broken TCP implementation. Therefore, the data by Shadowserver is used as the starting point for identifying IP addresses, and thus end-users, which are analysed within this research.

In chapter 4.2 the process of analysing the Shadowserver data is elaborated on.

4.1.2. Network scans

Information on vulnerable devices can be mapped by conducting three different scans. The following sections discuss the three scans in the chronological order in which they were conducted.

- **NMAP** NMAP is used for identifying open ports, services running on these ports and other information on the device of interest (Gordon Lyon, 2008). Scanning the IP addresses reported by Shadowserver with NMAP allows this research to identify the characteristics of most of the vulnerable devices.
- **Landing page analysis** To complement the information from the NMAP scans, the URL landing pages of the IP addresses (accessible via port 80) are analysed. When a device is accessible from the internet, the landing page of devices can provide insights on the type of device and brand name. Based on this data the devices could be further characterised.
- **Manual scanning** As the identified and analysed IP addresses are solely based on the reports by Shadowserver on vulnerable middleboxes, an additional scan is necessary to validate the

findings. The analysis of the data collected via NMAP and the URL landing pages indicates that not only middleboxes but other IoT devices can be the cause of the observed amplification. To validate this and find proof that these devices are not solely middleboxes, the IP addresses are manually scanned using a python script that sends custom TCP packets. This python script is based on the same scanning method as is used by Shadowserver except for one key difference. The python script sends packets without a forbidden URL, but instead an URL which should not be blocked by a firewall, IDS or another type of middlebox. As packets pass numerous hosts when travelling from the source to the destination, a middlebox in this path could block the packet and respond with a block page, resulting in an amplification factor (Bock et al., 2021). If the forbidden URL is replaced with a normal URL, it is assumed that the packets will no longer be considered interesting for a middlebox, and therefore, they should not respond with a significant blockage to the packet. The data returned by the vulnerable IP addresses is captured using Wireshark and analysed. This data validates which device is the cause for the reported amplification based on the HTML pages that are captured after sending the custom TCP packets.

This manual scanning has a disadvantage as it can only be used to validate the characteristics of devices which are still vulnerable. Devices which are not online anymore or devices where the broken TCP protocol is fixed for some reason will only respond to the custom TCP packets with an SYN-ACK packet instead of HTML data as is seen with vulnerable devices with broken TCP protocols.

The python script used can be found in appendix B.

4.1.3. End-users and manufacturer interview

End-users and a manufacturer involved in the issue of TCP reflective amplification within an ISPs network are interviewed. The goal is to examine the feasibility of using vulnerability notifications by an ISP by identifying end-user updating behaviour, their characteristics, notification perceptions, and actions. Besides, the findings of the interviews are used for creating an understanding of the role of these stakeholders in remediating the vulnerability.

All these interviewees are contacted on behalf of KPN in cooperation with the Delft University of Technology. KPN is able to link the contact information of actual customers to the IP addresses reported by Shadowserver. Interviewees are first sent an email, warning the end-users about the vulnerability with possible remediation steps. This email is sent via the KPN abuse channels, imitating the normal procedure of the ISP. Afterwards, every customer who received an email was contacted and asked for participating in this study via an interview. A detailed overview of the interview setup, notification protocol and interview protocol is discussed below.

4.2. Quantitative data analysis

As explained above, the daily data files by Shadowserver on vulnerable middleboxes serve as the starting point of this research. The format of these CSV files is shown in table 1. Each row corresponds to a unique IP address and includes the corresponding amplification factor, the scanning method used, and the date it was discovered. The raw data files provided by Shadowserver contain additional columns of information, however, these were not utilised for analysis in this research as they did not provide relevant information for the purposes of this thesis. An overview of these columns can be found in Shadowserver (2022).

Table 1: Example daily Shadowserver data

timestamp	hostname	amplification	method
1-1-2023 00:00:00	[IP ADDRESS 1]	100	SYN/GET
1-1-2023 00:00:00	[IP ADDRESS 2]	1000	SYN+ACK:PSH
...	...	...	...

For this research, all the daily files are combined and processed using Python. The output file includes

information for each unique IP address. This information includes the date the address was first and last observed in the data set, as well as the number of times it was recorded in the daily feeds between April 25th, 2022 and February 19th, 2023. Additionally, the highest amplification factor observed for each IP address is reported. Based on the responses received from scanning the amplification is calculated by dividing the received payload by the sent payload. Finally, the ratio between the time the address was first observed and the time it was last observed is also included in the file. The output of this analysis can be found in appendix F.

4.3. End-user interview setup

The end-user interviewees for this research are selected and contacted through the means of an interview setup. This section elaborates upon the process of target group identification, the vulnerability notification, the notification protocol and the end-user interviews. An overview of this process is visualised in figure 11.

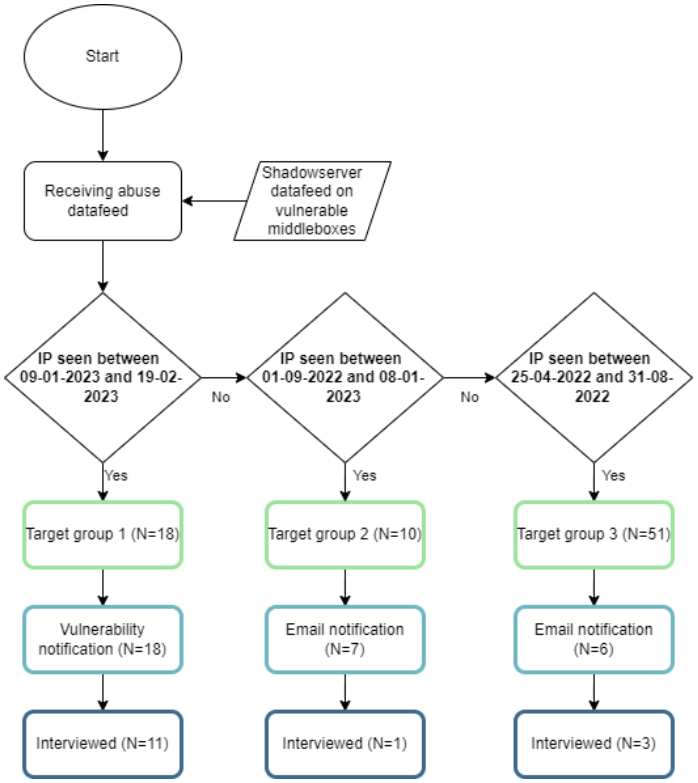

Figure 11: Overview end-user interview setup

4.3.1. Population of interest

The population of interest for this research are KPN customers who have a consumer account and own a vulnerable middlebox or IoT device which can be abused for TCP reflective amplification attacks. For this reason, KPN customers who are from the business or wholesale market are excluded. During the time period analysed in this research, a total of 2141 unique KPN IP addresses were found to have a vulnerable middlebox or IoT devices. The distribution of these IP addresses over the two markets is shown in table 2.

Table 2: Number of unique IP addresses per market type

Market type	Number of unique IP addresses
Consumer market (CM)	79
Business market (BM)	2062

Although IP addresses in the BM comprise 96.3% of the total population, they are not included in this research due to practical constraints. This is primarily because of the limited amount of time and resources available for the study.

The data provided by Shadowserver identifies 79 unique consumer IP addresses that have been reported between April 25th, 2022 and February 19th, 2023. All of these IP addresses are considered for analysis purposes in this thesis. The IP addresses can generally be divided into two distinct groups: those that were considered vulnerable during the research period, and those that were considered remediated during this research period. However, the assumed remediated IP addresses are further divided into two groups. This decision was made because in August 2022, reputable vendors of firewalls released patches fixing the vulnerability, and as a result, a large portion of the vulnerable IP addresses was no longer reported on by Shadowserver (see chapter 5). Therefore, the group of assumed remediated devices is split into two groups: one group with IP addresses last seen in August or before, and another group with IP addresses last seen between August and the research period.

This research identifies three different target groups based on the aforementioned criteria. The first group consists of IP addresses that were considered vulnerable during the research period. The second group is comprised of IP addresses that were considered remediated between August 2022 and the research period. Finally, the third group includes IP addresses that were considered remediated before August 2022. The specifics of each group will be discussed in greater detail below.

Target group 1: Recent vulnerable IP addresses The first target group consists of the IP addresses reported during the research period. As these IP addresses are vulnerable to TCP reflective amplification during the research period, target group 1 is considered the main population of interest. This period spanned five weeks, from January 16th to February 19th 2023. To identify vulnerable IP addresses, the research considered any IP address that appeared in the abuse data within the last seven days as vulnerable at the time of analysis. As a result, all IP addresses from January 9th to January 15th were also included in the first target group. In total, Shadowserver reported 18 IP addresses during this time frame, all labelled as target group 1.

Target group 2: Semi-recent vulnerable IP addresses The second target group consist of IP addresses that were last observed in the abuse data between the 1st of September 2022 to January 8th 2023. In total, Shadowserver reported 10 different IP addresses which were last seen during this time frame. All of these IP addresses were labelled as target group 2.

Target group 3: Historical vulnerable IP addresses The third, and last target group consist of IP addresses that were last observed in the abuse data between 25th of April 2022, when Shadowserver started reporting on this vulnerability and the first of September 2022. In total, Shadowserver reported 51 different IP addresses which were last seen during this time frame. All of these IP addresses were labelled as target group 3.

4.3.2. Notification protocol

For this thesis, a total of 31 customers were contacted. 18 of these were from target group 1, 7 from target group 2 and 6 were from target group 3. Depending on the target group, each of the customers received an email with either a vulnerability notification or an email informing the customer about the vulnerability and the research. This differentiation was made because customers in TG1 were deemed to have vulnerable devices during this research, whereas devices for TG2/TG3 were deemed remediated. After this, the contacted customers were called for participating in the interview. Below the process for sending the email and contacting customers is explained per target group. This explanation is combined for TG2 and TG3 as this process was identical.

Target group 1 notification protocol During the period of five weeks from January 16, 2023, to February 19, 2023, abuse data provided by Shadowserver is analysed for new IP addresses on working days. As this research considers IP addresses which are reported on a week before the start of the research period as vulnerable, all the IP addresses between January 9th to February 19th are included in the analysis and were contacted.

All of the corresponding customers, which can be linked to the identified IP addresses in TG1 are sent vulnerability notifications over email. This vulnerability notification will be elaborated on in chapter 4.3.3. 14 customers were contacted on January 17th, 2 on February 1st and 2 on February 13th.

Target group 2 & 3 notification protocol A limited amount of customers, which are linked to the corresponding IP addresses in TG2 and TG3, were sent an email on January 25th and February 1st respectively. In contrast to target group 1, where all customers were contacted, only a subset of the customers from target groups 2 and 3 were contacted. This was done as TG1 served as the main population of interest. For target group 2, seven respondents were randomly selected from the group. For target group 3, a total number of six customers were randomly selected. As the customers in these target groups are considered remediated, these customers did not receive a vulnerability notification as was received by TG1. This email received by these customers only explained the vulnerability in general and stated that he/she will be called for research purposes. For more information on the email sent, see chapter 4.3.3.

4.3.3. Email notifications

Depending on the target group, KPN customers were sent a notification over email before they were contacted for the interviews. Target group 1 received a vulnerability notification while target groups 2 and 3 received an email notification. This section elaborates on the two different emails that were crafted for this research.

Target group 1: vulnerability notification

Customers in target group 1 received a vulnerability notification providing them with information about the vulnerability, the type of devices which presumably is vulnerable and how to remediate the vulnerability. This vulnerability notification can be found in figure 19 and figure 20 in appendix C. During the construction of the email, various decisions are made in regard to the content. Although the notification is not part of the data collected through this research, it can have a substantial impact on the willingness of customers to participate and their thinking process around handling such emails and their vulnerable devices.

To make sure that it the read and understood by as many end-users as possible, the text and content of the notification are reviewed by three abuse experts from the KPN abuse desk, who have experience in communicating such content. Due to technical constraints, the text was only written in Dutch and not in English. However, it can be noted that the communication between KPN and end-users is usually only written in Dutch, and did not result in difficulties in the past as explained by KPN abuse desk employees.

The email begins by briefly describing the issue and what devices are likely to be causing the problem. First, devices are named that are part of people's internet infrastructure. The email states explicitly that firewalls or routers/modems are the likely cause of the issue. The first reason for naming these devices first is the fact that Shadowserver scans for vulnerable middleboxes. As explained by Bock et al. (2021), firewalls are often broken devices making TCP reflective amplification possible. Moreover, routers and modems often have firewall settings that can be utilised. Therefore these are mentioned explicitly as well. Additionally, routers are found to often have vulnerabilities, making them a suspect for the issue of TCP reflective amplification as well (Katwala, 2019; Wang et al., 2012). It is not expected that ”normal” consumers would have a more intricate IT infrastructure in their homes beyond a router and a modem. As a result, the email does not explicitly mention the term ”middlebox” or other examples of middleboxes besides firewalls.

Secondly, a list of 4 different consumer IoT devices has been mentioned. These are energy monitoring systems, alarm systems, cameras and building management systems. These are mentioned as these were the four different IoT devices that were seen during the initial exploration of the vulnerable devices.

However, after validating these initial findings it was found that the IoT devices are indeed the cause for the observed amplification by the Shadowserver scans. These findings are presented in chapter 5.

Last, the email stated that the issue can be remediated by updating the devices named in the email. During the research period, when customers received the vulnerability notification it was expected that updating does not have the desired outcome for all devices. Even though updating devices is a suggested way of reminding the device (Pal, 2022). The main reason for this is the assumption that not all manufacturers are aware of the broken TCP protocols in their devices and consequently have not yet released updates to fix the vulnerability. This assumption was later verified by a manufacturer who is responsible for the majority of broken devices in TG1. This finding is described in chapter 6.2. Nevertheless, it is decided to state updating the device as the step that needs to be performed for remediating the vulnerability. As the goal of this thesis is to explore the issue within the network of an ISP to ultimately help ISP's with remediating the vulnerability, by doing so the behaviour of end-users receiving such notifications can still be observed.

When formulating the content, findings from the literature study, as stated in chapter 2.3, were taken into account. The notification needs to be comprehensive yet easy to understand for non-technical people. The difficulty in crafting the notification was the number of different devices which are likely causing the problem. That is why the notification was generic in nature and loosely stated that "updating smart devices" is the needed action for remediating the vulnerability. The advantage of this is the fact that it is easy to understand for everyone, yet it lacks specific steps that possibly obstruct people from performing the intended actions. Even though it is likely for the notification to lose effectiveness, this generic content is chosen. To provide more (technical) information on the vulnerability, a link was included to the website of the ISP with additional information.

The link to the ISP's website had a dual purpose. It was discovered in the literature that mistrust can have a significant impact on the efficacy of vulnerability notifications. By including a link to KPN's website, individuals could independently verify the authenticity of the email, which may have increased customer trust. The website is included in appendix C in figure 23.

Target group 2 & 3: email notification

Contacted customers from target groups 2 and 3 received another email than the customers from target group 1. Compared to TG1, for TG2 and TG3 only a subset of the population received an email. The customers from these target groups were randomly selected. The main purpose of the email was to announce that they will be contacted for this research. Even though the email explained the vulnerability, it clearly states that at the time of sending the email the vulnerability seems remediated, thus no actions of the customer are needed. The notification can be seen in figure 21 and 22 in appendix C.

Similar, to the vulnerability notification sent to TG1, a link to the website of KPN was added with extra information on the research and the vulnerability. This was done to gain extra trust of the receiver. Furthermore, to gain trust the name of the researcher was stated in the email which gives the opportunity for the customer to link the phone call to the received email, possibly increasing the trust of the customer. The website can be found in figure 23 in appendix C.

4.4. End-user interviews

After the KPN customers received either a vulnerability notification or an email notification, they were called for interview purposes. The aim of the interviews is to understand four key aspects of the end-users. First, the behaviour of end-users after receiving a vulnerability notification. Second, their behaviour in regard to updating smart devices within the household. Third, end-users perception of a vulnerability notification by the ISP. Last, the characteristics of the end-users. Additionally, the interviews are used to verify the findings of the vulnerable devices, as described in chapter 5.

The end-users interviews are conducted in a semi-structured matter. Semi-structured interviews are suitable for studies with an exploratory character (Bougie et al., 2019). This type of interview allows this research to both compare data collected among the end-users as well as to collect new insights (Horton et al., 2004). Furthermore, conducting semi-structured interviews to identify end-user behaviour and

perception in the field of cyber-security has been conducted by various studies before (Anell et al., 2020; Rodriguez et al., 2022; Tabassum et al., 2019).

4.4.1. Interview protocol

The interview protocol consists of two parts: an introduction part and the content part. The interviews are conducted over the phone and the interviewees have no knowledge about the interview, the research or the vulnerability prior to receiving a vulnerability notification or email notification. That is why a comprehensive introduction is needed before asking the interview questions in a semi-structured manner.

Preferably, this interview protocol goes through a pilot phase where the interview questions are tested in order to make informed adjustments to the interview protocol and improve the quality of data collection (Kallio et al., 2016). However, for this research, no pilot study was performed. The reason for this was the small population size of vulnerable IP addresses. To ensure the quality of the interview several steps have been taken before conducting the interviews. First, the pilot studies and interview protocols from similar studies conducted at KPN were taken into account when formulating the protocol and questions (Altena, 2018; Bouwmeester et al., 2021; Rodriguez et al., 2022; Rodríguez et al., 2021). Even though, the content of these studies differs from the content covered in this thesis, the experiences from this previous research could be used as a basis for formulating the interview protocol for this study. Furthermore, the interview protocol was reviewed by TU Delft and KPN experts who have experience with conducting similar studies and end-user interviews respectively.

As there are different target groups, there are two interview protocols. One for target group 1 and one for target groups 2 and 3. The two different interview protocols are discussed below.

Interview protocol TG1

The introduction part can be found figure 24 in appendix D. The interview starts with introducing the interviewer and the research. Besides, there are three things that need to be clarified before being able to start the second part of the interview and ask the questions. First, it is important to make sure that the person calling is the same person who is responsible for the security of the computer and the other (smart) devices in the household. Second, it is necessary to ask whether the customers received and read the vulnerability notification as otherwise the interview questions do not make sense. Lastly, respondents were asked for informed consent over the phone which is required for conducting an interview for research purposes. When these three criteria are met, the interview could proceed to the second part.

The second part of the interview protocol, containing the interview questions, can be seen in figure 25 in appendix D. These questions are divided into 5 topics which were covered during the interview. First, interviewees are asked about the (smart) devices in their homes, whether they were able to locate the vulnerable device and if they bought the devices by themselves. As there is a large heterogeneity among the smart devices with broken TCP protocols it can be hard for end-users to pinpoint the vulnerable device. Additionally, interviewees are asked whether they purchased the devices themselves to understand their level of responsibility towards the device. The responsibility for updating a device installed by an external party may lie with that party, relieving the user of the obligation to update it. That is why, understanding the smart device landscape in households, including the number of devices and perceived responsibility, as well as the ability of end-users to identify vulnerable devices is crucial.

Second, interviewees are asked about their performed actions after receiving the vulnerability notifications. By doing so, this research tries to identify the thinking process of the interviewees when receiving a notification about the issue of vulnerable middleboxes or smart devices with broken TCP protocols. In addition, interviewees are asked if they think they had enough information for performing the desired actions, this will provide this research with possible suggestions for future vulnerability notifications.

Third, interviewees are asked about the updating behaviour in regard to their smart devices. In this research, updating the vulnerable device is the suggested remediation action, and it is the desired action that interviewees are expected to perform after reading the vulnerability notification. Therefore,

this aspect of the interview is essential for studying the effectiveness of the suggested remediation action. The interviewees are asked how they perform the updates on their smart devices and their incentives for (not) updating. This will provide insights into the updating behaviour of the interviewees, which can be used to assess the feasibility of vulnerability notifications. Furthermore, the interviewees are also directly questioned about their opinion on whether vulnerability notifications effectively lead to the desired action of updating suggested smart devices.

Fourth, the interviewees are asked about their perception of receiving a vulnerability notification and the service of an ISP of notifying its customers about vulnerabilities. As vulnerability notifications aim to encourage end-users to perform a certain behaviour it is interesting to map the perception of the end-users. Moreover, the interviewees are asked if they have prior experience with security incidents or vulnerability notifications sent by the ISP.

Last, the interviewees are asked to provide some demographic information. This will provide insights into the characteristics of the end-users and how these differ among the different target groups. Besides demographic information, the interviewees are asked whether they use the internet connection for business or private purposes. This research focuses on the consumer market instead of on the BM or WM. However, previous similar research conducted at KPN showed that there are cases of businesses using a consumer internet connection. This is also confirmed by KPN abuse desk employees, as a consumer internet subscription is cheaper for small businesses that do not need complex IT infrastructure for their business activities.

The interview is finalised by asking the interviewee if he or she has any questions or comments. As the phone call is performed in the name of KPN, it is expected that interviewees have questions about the vulnerability and the security of their devices or internet connection. This allows the respondents to seek assistance or clarification regarding the matter at hand.

Interview protocol TG2 & TG3

The interview protocol for TG2 and TG3 is largely the same as for TG1, however, there are some differences which will be discussed in this section. These differences are mainly there because these two target groups have no vulnerable device during the research period and subsequently these customers did not receive a vulnerability notification. As the email notification that was sent, primarily serves as an announcement for the interview, the first part of the interview protocol does not inquire whether the interviewee received or read the notification. An additional distinction is that no inquiries are made regarding the actions taken after receiving the vulnerability notification, as no notification is sent to these interviewees.

Part 1 and part 2 of the interview protocol for TG2 and TG3 can be seen in figure 26 and figure 27 respectively, both in appendix D.

Interview (questions) in relation to the COM-B model

This research applies to the COM-B model to analyse the findings from the interviews. Therefore, it is important to understand how the interview questions address the three interacting components as discussed in chapter 3.3. The capability component, in terms of this research, is the ability of end-users to have the physical and psychological ability to perform the actions as intended by the vulnerability notification. The study does not directly address physical ability as the intended actions based on the vulnerability notification only require a vulnerable device and a smartphone or computer for updating the device. Prior to sending the notification email, the presence of the device is verified, although this is confirmed during the interview process. Additionally, it is assumed that all participants in the study possess a smartphone or computer, and this is not explicitly queried. However, in instances where the interviewer suspects this may not be the case, clarification is sought directly from the participant during the interview. Furthermore, by identifying the thought process of end-users to perform certain actions, this research is able to explore the comprehension and reasoning of end-users, thus addressing their psychological ability.

For this research, the opportunity component can be described as all the external factors of the end-user that make the intended actions possible for end-users to perform. In other words, external factors may influence the end-user in performing or not performing a certain behaviour. The physical opportu-

nity is addressed by asking whether the respondent received the vulnerability notification and by asking if the end-users have enough information (and what information might be missing). The social opportunity describes all the factors from an end-user environment which may have an effect on the thinking process of the individual when receiving the notification and performing certain behaviour afterwards. As this may be ambiguous and hard to measure, this research tries to address this by performing a semi-structured interview, allowing the interviewer to go in-depth into the responses provided by the interviewees.

The motivation component refers to the cognitive and affective processes that activate and direct specific behaviours in response to the vulnerability notification. This study distinguishes between two types of cognitive processes, namely reflective and automatic, which are involved in evaluating and planning actions after reading the notification and in perceiving the notification as a signal of potential harm or another type of emotion. These cognitive processes are fundamental to understanding how end-users respond to vulnerability notifications and how they ultimately decide whether to take action to remediate the vulnerability. Therefore, in the interview, there are questions included about end-user perceptions of receiving vulnerability notifications.

4.4.2. Interview results

In total 31 end-users are contacted of which 15 are interviewed, a response rate of 48.4 per cent. However, it must be noted here that one of the contacted end-users was actually the manufacturer of the energy monitor, often found as vulnerable (see chapter 5). Therefore, this interviewee is not included in the end-user interviews but is separately interviewed. For more information see the following subchapter 4.5. On average the interviews took 11 minutes and 20 seconds. An overview of the contacted customers is presented below in table 3.

Table 3: Contacted customers for the interviews

	Number of customers
Participated	15
Not participated	8
Did not pick up	6
Number not in use	1
Manufacturer	1
Total	31

Ideally, one would conduct interviews until no new themes emerged and information saturation was reached (Guest et al., 2006). However, due to the small population size all customers in TG1 are contacted for interviews. Because the main focus of the interviews was on TG1, from the TG2 and TG3 a subset was selected to reach the desired level of information saturation and no new themes emerged.

4.4.3. Processing interview results

All interviews, except one, were recorded. One interviewee (respondent 8) expressed his/her desire to not record the interview. The recordings are transcribed after all interviews had taken place and irrelevant small talk was left out of the transcription. Using these transcriptions, the qualitative data was analysed. For analysing the data Atlas.ti software was used.

A thematic analysis is performed to identify recurring themes and to structure the patterns across the data set. To do this, a six-phase process for data analysis, coding and theme development is used which is proposed by Braun et al. (2021). The phases are as follows: first, data familiarisation; second, systematic data coding; third generating initial themes from the codes; fourth, developing and reviewing themes; fifth, refining, defining and naming themes and the last phase writing the report. This process resulted in the identification of the following themes shown in table 4.

Table 4: Identified themes from end-user interviews

Themes	Description
End-user characteristics	Descriptive data on age, gender, type of customer and technical ability
Device identification	How end-users identify the vulnerable device
Performed actions	Which actions end-users have performed after reading the vulnerability notification
Update behaviour	Which practices end-users use for updating their (smart) devices
Incentives for (not) updating	Different reasoning for end-users to (not) update their (smart) devices
Perception of vulnerability notification	Opinions and emotions of end-users after receiving the vulnerability notification and during the interviews
Suggestions for vulnerability notifications	The different types of information end-users perceive to need for a vulnerability notification to be effective

4.5. Manufacturer interview

This study also included an interview with the manufacturer of vulnerable energy monitors. These devices were found to be responsible for the observed amplification reported for the majority of the IP addresses in TG1. As detailed in Chapter 5, for 11 out of the 18 customers in TG1 this particular energy monitor is identified as the vulnerable device causing the amplification and thus makes TCP reflective amplification attacks possible.

As previously noted, the manufacturer received a vulnerability notification as his/her IP address is in TG1. In response to the notification, the manufacturer expressed responsibility for the vulnerable energy monitors. As a result, the manufacturer was invited to participate in an interview as a representative of the manufacturing company, rather than as an end-user, despite being classified in TG1 and using a consumer-grade internet connection from KPN.

The objective of the manufacturer interview is to explore the role of the manufacturer in the problem of IoT devices with broken TCP protocols in the network of KPN. Therefore, the data retrieved from the interview can be seen as an addition to the data obtained during the end-user interviews. By gaining insights into the role of the manufacturer in the issue of interest, this interview may provide new perspectives on the themes and findings from the end-user interviews.

Interview protocol and data processing

The interview is conducted in a semi-structured fashion by phone and took 26 minutes and 34 seconds. The interview is recorded and transcribed afterwards. The interview questions are displayed in appendix E. As only one manufacturer was interviewed for this research, no systematic theme analysis is performed to identify themes and structure patterns across the data set. Instead, the study focuses on the key takeaways that were noted during the interview which are relevant in relation to the findings from the end-user interviews.

The purpose of the first question is to ascertain the manufacturer's knowledge and experience regarding the vulnerability of TCP reflective amplification in the context of the energy monitor devices. The objective was to determine the manufacturer's awareness of the potential risks associated with these devices and to evaluate whether any corrective measures had been taken or could have been taken. Obviously, it is challenging to mitigate a vulnerability when one is not aware of its existence. Furthermore, the questionnaire also aimed to explore the manufacturer's response to the vulnerability notification, including their actions and thoughts, as well as their standard mitigation protocols for vulnerabilities in the energy monitor devices. Lastly, the interview aimed to ascertain the methods employed by the manufacturer to communicate security updates to customers and encourage them

to implement these updates. This analysis was conducted to compare the manufacturer's responses to the experiences of end-users and to examine the manufacturer-end-user interaction in regard to security updates for these devices.

4.6. Limitations of the method

While this chapter outlines how the method can offer insights into the primary research question and the sub-questions, it is crucial to note that there are various limitations. This section will elaborate upon the limitations of the method employed in this study.

First, it must be noted that this study is exploratory in nature, and as such, it does not seek to quantify the findings or make generalisable conclusions about a larger population. Rather than focusing on generalisability, the aim is to achieve a rich understanding of the topic. The findings, emerging themes, and discussions arising from the analysis should be tested in a subsequent quantitative study to assess their wider applicability and generalisability.

An additional constraint of this method is the reliance on Shadowserver's data on which the following steps in the method are built. However, it is important to understand that Shadowserver only scans via port 80 with a single forbidden URL. As described by Shadowserver (2022), scans using different URLs' or via other ports could result in a different number of identified IP addresses and observed amplification factors.

Furthermore, Shadowserver only performs one scan per day at random times per IP address. Therefore, it could be that vulnerable devices are not connected to the internet at all times, and thus do not respond to the scans of Shadowserver. For this research, IP addresses which are last seen in the abuse data seven days before the research period are considered remediated. Nevertheless, it could be the case that a device is not remediated at all but simply is turned off every time Shadowserver performs its scans.

The main data sources used for this research are end-users of vulnerable devices which provided data by the means of interviews. Therefore, it can be assumed that interviewees give answers that are biased, and thus give answers which differentiate from reality. The interviewees may have forgotten the actions they performed or unremembered them, leading to inaccuracies in their responses. Furthermore, respondents may provide answers that they believe are desired, particularly since the notification emphasises the importance of updating smart devices, and the interviewee is questioned about their update behaviour during the interview. To overcome these interview biases, the interview is framed in such a way that it can be seen as an effort to help the interviewees, with the aim of encouraging them to share their thoughts and update behaviours more freely.

Another limitation to the limited sample size of the interviewees. There are two reasons for this limited amount of conducted interviews. First, the main population of interest, the population of IP addresses vulnerable during the research period (TG1), only consisted of 18 in total. Second, a limited amount of resources and time was available for this research. While it could be harder to identify common themes and patterns among the data set obtained via the interviews, a response rate of 48.8 per cent is achieved, which is similar to previous research studying end-users at KPN (Altena, 2018; Bouwmeester et al., 2021; Rodríguez et al., 2021). Subsequently, due to the small size of the population of interest, a pilot study was not conducted to test the interview protocol and questions. The choice to omit a pilot study is discussed in chapter 4.4.1. This could have impacted the quality of the interviews and the resulting data, making it a limitation of the method.

The last limitation of the method which can be mentioned is that only one vulnerability notification is sent to the interviewees in TG1. However, it can be expected that multiple notifications could affect the update behaviour, actions and perception of end-users. This factor is described by previous research that found that end-users sometimes require multiple notifications before they act (Rodriguez et al., 2022). The effect of one or multiple vulnerability notifications is not discussed as the effectiveness of vulnerability notification is not in the scope of this research.

4.7. Ethical considerations and data management

For this research personal data is used. Therefore, prior to the start of the study, the Human Research Ethics Committee of TU Delft reviewed and approved the research. Furthermore, the data processes are in line with the General Data Protection Regulation (GDPR) and KPN privacy statement. A senior Abuse desk employee approved the necessary steps for this research which involved handling personal data, such as sending the vulnerability/email notifications and conducting the interviews.

As the interviews are performed over the phone, all participants were asked for informed consent for participating in the research during the interview. Furthermore, the interviewees were asked to give approval for recording the interview. Contact details of customers are only looked up right before either emailing or calling the customer, this data is not collected for research purposes. All data collected during the interviews, including the recordings, are stored on a KPN laptop and TU Delft Microsoft OneDrive. This data is solely used for analysing and processing the results of this research. All results published in this report are anonymised and no personal data is shared which can be used for identifying the participants of this research. All interview recordings are destroyed after processing.

5

Vulnerable devices in the KPN network

This chapter presents the outcomes derived from analysing the abuse data provided by Shadow Server and performing IP scans to identify the characteristics of vulnerable devices that enable TCP reflective amplification attacks within KPN's network. First, the results from the abuse data analysis are presented (§5.1). Second, the results from the vulnerable device characterisation are shown for each of the three target groups (§5.2). Last, the results in regard to the identified device types is discussed (§5.3).

The next chapter will discuss the findings from the end-users and manufacturer interviews. However, for the purpose of elaboration on the results in this chapter, some findings from the interviews will be referred to. All other findings from the interviews will be presented in chapter 6.

5.1. Vulnerable devices in the consumer market

First, an overview of the vulnerable IP addresses is analysed within the consumer market of KPN. In total between the 25th of April 2022 and the 19th of February 2023 79 unique IP addresses have been identified which are vulnerable to being used for launching TCP reflective amplification attacks. This is visualised below in figure 12. This figure shows the number of uniquely identified CM IP addresses per week.

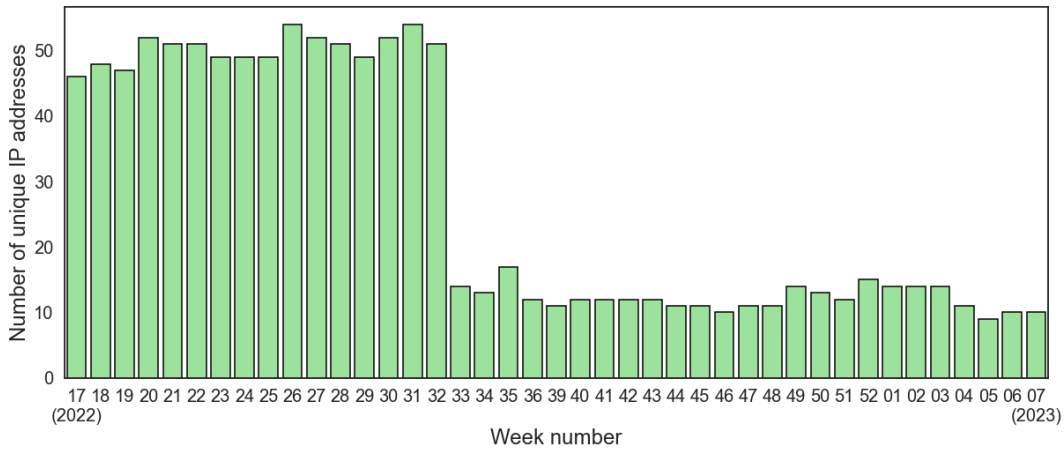

Figure 12: Number of unique KPN CM IP addresses per week

There are three interesting things that can be observed within this figure 12. First, there is a large drop observed in the number of unique IP addresses between week 32 and week 33 in 2022. This can be explained by the fact that large vendors of firewalls released their patches for the vulnerability as

described by Bock et al. (2021). The release of the patches by these large vendors is around the same time that the number of unique IP addresses drops from around 50 per week to around 15 per week - a decline of 70 per cent (Fortinet, 2022; Palo Alto Networks, 2022).

Another notable observation is the minor decrease in reported IP addresses during weeks 3, 4, and 5 of 2023, which correlates with the period when the end-users received the vulnerability and email notifications and were contacted for partaking in an interview. It is possible that some individuals may have taken action to rectify the vulnerability upon receiving the notification however this was not expected based on the content of the notification. For example, during the interview with the manufacturer, it was revealed that he/she had closed port 80, resulting in its lack of response to the scans executed by Shadowserver. However, the reduction in activity is not substantial. It was found that updating during the research period did not always result in the desired outcome of resolving the vulnerability. Therefore, it was expected to not observe a drop in the number of unique IP addresses during the weeks of the research period.

The final observation that can be made from this figure is that the count of unique IP addresses differs for nearly every month when compared to the previous or subsequent month. In addition to the two drops in the number of unique IP addresses, as previously discussed, it is possible that the variation in the number of unique IP addresses could be attributed to devices with broken TCP protocols being turned off during the daily scans. Not all IoT devices require a constant internet connection. For instance, consider a washing machine that only needs to be connected to the internet when it is in use to communicate when the wash cycle is complete. Consequently, while the device itself may not be detected by the scan, the vulnerability still remains present.

5.1.1. Descriptive statistics insights

In figure 13 the maximum amplification factors are shown for each device type for all target groups combined. The graph has been segmented into 5 amplification ranges, depicting the highest observed amplification factors for each device type and the corresponding number of devices generating such levels of amplification. Note, *NA* are the unidentified devices.

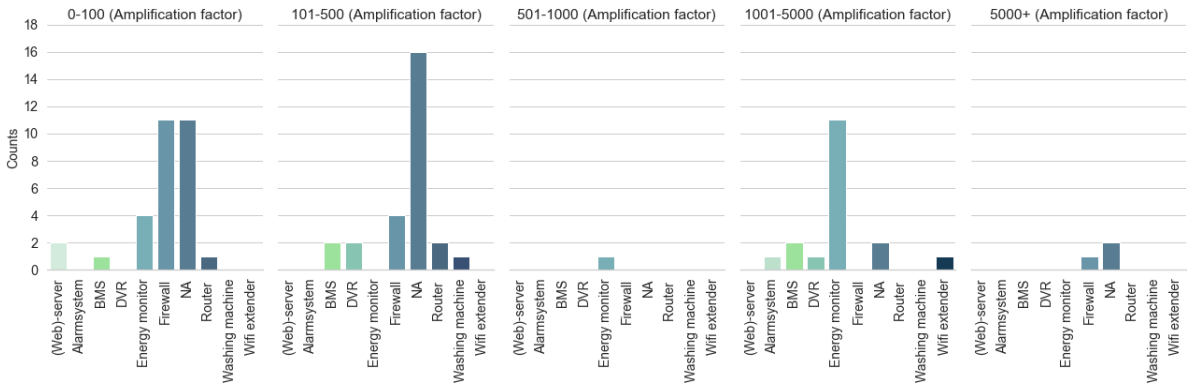

Figure 13: Distribution of maximum amplification factors

An overview of the descriptive statistics can be found in appendix F. Based on the descriptive data a number of interesting remarks can be mentioned in regard to the data on the vulnerable IP addresses in the consumer market.

The observed amplification factors, as depicted in Figure 13, indicate that the majority of the identified devices generate amplification factors of up to 500. It is worth noting that the 1001-5000 range exhibits a peak in energy monitors, with most of them belonging to TG1 along with an alarm system, WiFi extender, and BMS found within this same range. This observation demonstrates that, at the time of writing, there are multiple vulnerable devices within KPN's consumer network that could be exploited for significant TCP reflection amplification with amplification factors exceeding 1000. In addition, three devices exhibited amplification factors over 5000, but they can be considered outliers since they were substantially larger than the others, with the firewall in this range even reaching an amplification factor

over 150000.

Secondly, the average maximum amplification factor that was found is 2801. This is considerably higher than average amplification factors found in commonly used UDP-based protocols (Bjerre et al., 2022). The average amplification factors differ between the different target groups. The average maximum amplification factor in TG1 and TG2 is 1111 and 1121 respectively while this number is for TG3 3727. This difference between TG3 and the other target groups can be explained on the basis of two reasons. First, the composition of the type of devices is in TG1 and TG2 quite similar, as is explained below, consisting of almost exclusively consumer IoT devices. TG3 consists mostly of firewalls, which are likely causing the problem and appear to generate another amplification factor. The other reason is that there are three large outliers in the data resulting in a significant increase in the average maximum amplification factor. Without these outliers the average is ”only” 167, nevertheless, this is still high.

Third, it is noteworthy that the observed variation in amplification among the energy monitors, as seen in figure 13, is attributable to the availability of a password-protected mechanism. Despite the devices being identical, the responses elicited by the scanning query are contingent upon the presence or absence of a password. Specifically, in the absence of a password, the HTML response contains a greater amount of information, as it showcases the entire landing page of the device displaying all the different energy values that the device collects for monitoring purposes. The amplification factor is approximately 1600 as a result. In contrast, when a password is in place, the device generates a more concise landing page. This results in an amplification factor of approximately 71. This landing page displays only a login mechanism, in contrast to the complete landing page displayed in the absence of a password. As a result, the amplification factor produced by the device is substantially smaller.

Last, it is notable that a majority of the IP addresses are not reported on a daily basis in the time period between their first and last appearances in the data. On average, only 62.5 per cent of the days between a vulnerable IP’s first and last appearance in the abuse data, it is reported by Shadowserver. This may be partially explained by devices not being consistently turned on. However, some devices are expected to remain connected to the internet at all times. The IP scans revealed that devices do not respond to a scan every single time, thus sometimes the custom TCP packets need to be sent multiple times to create a response which causes the amplification. The reason behind the device’s unresponsiveness to the scans by Shadowserver, despite no alterations being made to them, remains unclear. However, it does explain why vulnerable devices are not reported on by Shadowserver every single day.

5.2. Types of vulnerable devices

This sub-chapter shows the different types of identified vulnerable devices in the consumer market of KPN. The findings are shown for each of the three target groups.

5.2.1. Target group 1

Based on the desk research and the interviews, as discussed in chapter 4, the different device types are identified which cause the amplification that makes TCP reflected amplification attacks possible. The distribution of these devices is visualised in figure 14.

Based on the information in figure 14, it can be seen that there is a wide variety of device types identified and that these all do not fit the description of a middlebox. Out of the 18 IP addresses in TG1, it was found that eleven times (61.1%) an energy monitor is the likely cause for the observed amplification at the IP address. All these 11 energy monitors are the same device in terms of brand and type and can be used for monitoring the energy usage in a household. The other types of vulnerable devices are less represented in TG1, with three building management systems (BMS) and only one alarm system, washing machine, WiFi extender and DVR. No similarities were found in the brand except for the energy monitors and two BMS’s.

By manual scanning, these findings could be validated and the vulnerable devices could be pinpointed. The energy monitor, alarm system, washing machine, WiFi extender and a BMS have been characterised based on this technique. However, this method was not able to validate the device types for all the IP addresses in TG1. A possible explanation for this could be that devices are turned off sometimes,

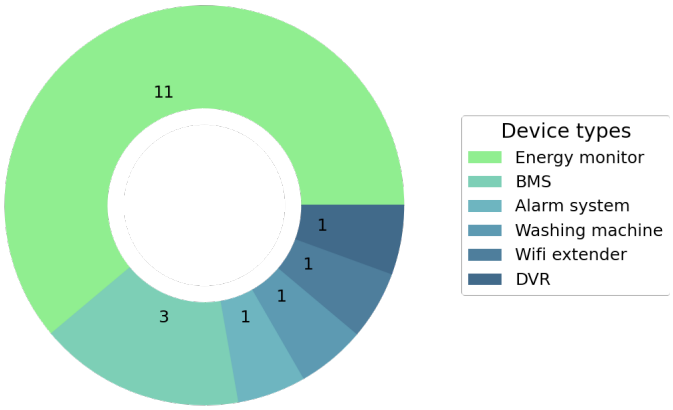

Figure 14: Distribution of vulnerable devices in TG1

and thus are not always connected to the internet. This was observed during one of the interviews with an end-user (respondent 8). The respondent turned on the washing machine during the interview, after which a response was created by manually scanning the IP address, validating that the washing machine was the device causing the problem. Therefore, it is assumed that the devices identified based on the NMAP scans, URL landing page analysis and interviews are the ones causing the problems but nevertheless, it can not be concluded with complete certainty.

5.2.2. Target group 2

The different vulnerable device types that were identified for the IP addresses in TG2 are shown in figure 15.

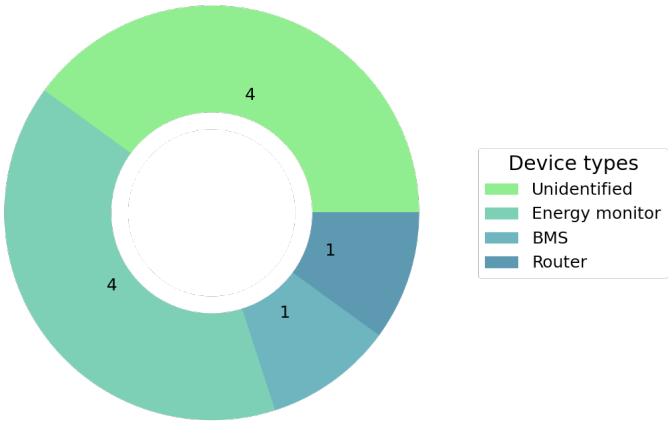

Figure 15: Distribution of vulnerable devices in TG2

Out of the 10 IP addresses, four times the energy monitor (also found in TG1) was identified as the probable cause of amplification. This is noteworthy as within the recent abuse reports, these IP addresses are no longer reported on, even though three of them were still accessible via the internet as of the last day of the research period. This lack of reporting may suggest that the malfunctioning TCP protocol has been fixed. However, during an interview with the manufacturer, it was revealed that no patch has been released for this particular issue. Furthermore, the only interviewee from TG2 mentioned that he/she had not made any changes to their device or internet connection around the last date the IP address was reported. The reason why these IP addresses no longer appear in recent abuse feeds is unclear.

Furthermore, for four IP addresses, it was not able to identify a device type. The last two IP addresses were identified as a BMS and a router. As explained before, it is not able to verify these findings as the IP addresses did not respond with an amplification anymore when manually scanning as these are

remediated.

5.2.3. Target group 3

The distribution of the different vulnerable device types that are identified for the IP addresses in TG3 is illustrated in figure 16.

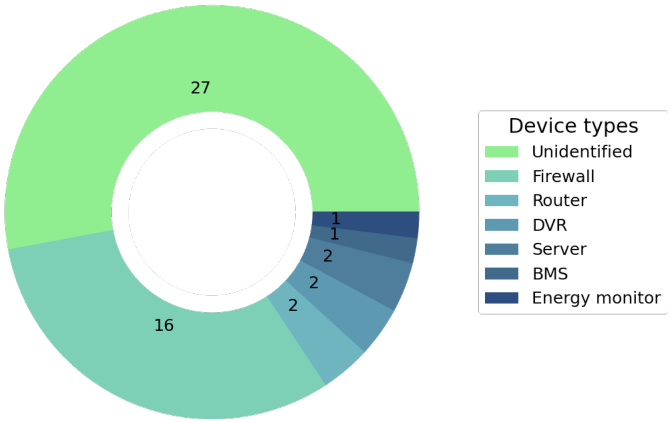

Figure 16: Distribution of vulnerable devices in TG3

From figure 16 it becomes apparent that for most (55,1%) of the IP addresses it was unable to identify the type of device. The second largest share of the IP address had a firewall identified as the likely cause for the observed amplification. As described above, around the beginning of August 2022 a large drop was observed in the number of weekly reported unique IP addresses which are vulnerable to TCP reflective amplification attacks. Around 35 IP addresses were reported on for the last time during this time period of which all were identified as firewalls or unidentified. Therefore, this suggests that the unidentified devices in this group of 35 IP addresses are also (mainly) firewalls. The three interviewees from TG3 confirmed this assumption during the interviews. However, this can not be concluded for the whole population as no data on this matter is available for analysis. Therefore, in figure 16 these IP addresses are labelled as ”Unidentified”.

The other types of identified devices are similar to what was observed in TG1 and TG2; consisting mainly of consumer IoT devices. Alike TG2, it can be seen that an energy monitor (same brand as the energy monitors in TG1 and TG2) was vulnerable but is somehow remediated for creating large amplification factors.

5.3. Results analysis

A large number of devices characterised in this chapter can not be characterised as middleboxes, even though Shadowserver reports on them as vulnerable middleboxes. This finding can be seen as quite surprising as it is not expected for other devices to respond to requests with an obscene URL in the request. Bock et al. (2021) conclude from their research that 82.9% of their identified amplifiers are actual middleboxes by looking at the traceroutes to the IP addresses.

For this research, the findings about the type of devices were validated by sending the custom TCP packets manually to observe what response is generated. This method was only suitable for validating the findings in regard to the IP addresses in TG1. Nevertheless, the data from TG1 already gives valuable insight; the scanning method for detecting vulnerable middleboxes does not exclusively detect these vulnerable middleboxes but also (consumer) IoT devices with broken TCP implementations. Based on this conclusion one may suggest that the findings for TG2 and TG3, which are not validated by manually sending TCP packets, are accurate solely on the analysis of the port scans and the URL landing pages.

When excluding the unidentified devices, it was found 16 of the 48 (33.3%) identified vulnerable devices are a firewall and thus can be described as middlebox. This number is substantially lower than the 82.9 per cent found by Bock et al. (2021). There are two possible explanations for this lower number.

First, as described above, the data suggests that a large part of the unidentified devices are firewalls and thus can be classified as middleboxes. Second, Bock et al. (2021) explain that a large part of the middleboxes that were identified are part of nation-state censorship infrastructure in countries like China and Iran. As there is no such infrastructure present in The Netherlands it can be expected that fewer middleboxes are detected when utilising the scanning method in the network of a Dutch ISP.

6

End-user analysis

In this chapter, the findings from the end-user and manufacturer interviews will be detailed. First, the results of the end-user interviews are presented (§6.1). This section is structured according to the main themes of this analysis, as shown in table 4. The structure looks as follows:

- End-user characteristics (§6.1.1);
- Device identification (§6.1.2);
- Performed actions (§6.1.3);
- Update behaviour (§6.1.4);
- Incentives for (not) updating (§6.1.4);
- Perception of vulnerability notifications (§6.1.5);
- Suggestions for vulnerability notifications (§6.1.5).

Second, the results from the manufacturer interview are presented (§6.2). Lastly, the end-user behaviour is analysed by linking the interview findings to the theoretical constructs of the COM-B model (§6.3).

6.1. End-users

In this sub-chapter, the findings of the interviews with KPN customers are discussed. As described in chapter 4, three distinct target groups were identified and were asked about their vulnerable devices and their security behaviour in order to explore the problem discussed in this thesis. The study participants in TG1 were interviewed regarding their response to a vulnerability notification. The interview aimed to understand their actions after receiving the notification, their smart device update behaviour, their perception of the notification and the KPN service. The study participants in the TG2 and TG3 groups were interviewed but were not asked about their actions in response to the vulnerability notification as they did not receive such notification.

6.1.1. End-users characteristics

During the interviews, five questions were asked to identify the characteristics of the end-users. This was done for all three target groups. First, the age of the interviewees was obtained. The average age of the interviewees is 60 years old for all three target groups combined. This is slightly higher compared to the average age of the KPN customer population, being 55 years old.

The average age for the first target group was 60 and for the third target group 70. Only one person (respondent 11) was interviewed from target group 2 who was 34 years of age, and thereby also the youngest interviewee. Overall, the average age is relatively high and not widespread. There are two outliers, respondent 11 has an age of 34 which is substantially smaller than the average age of

the interviewees and respondent 12 has an age of 86 which is substantially larger than the average age.

Second, the gender of the interviewees was asked. 12 of the 15 interviewees identified themselves as male and the other 3 (20%) identified themselves as female. This shows that males are over-represented in the subset KPN customers who were interviewed, compared to the gender distribution among the KPN customer population, being 55,64% men,41,38% women and 2,98% unknown.

Third, interviewees were asked about their perceived skills in computer security in general on a scale of 0 to 5. On average respondents perceive their computer security skills to be rather good with an average score of 3.9. The responses of most interviewees were in line with their own perceived computer security abilities. Multiple respondents even mentioned that they have a background in ICT or cyber security (respondents 1, 5, 6 and 14). The lowest score that was observed was a 3. Interestingly, some interviewees showed during the interviews that they perceive their ability to handle the security of their devices on a higher level than their responses suggest. For example, during the interview, respondent 2 argued the following when asked about the actions performed after reading the notification that was received: *'I must admit that normally I do not do a lot with these type of things [computer security], so I feel a bit foolish but I just do not know what to do otherwise'*. While this suggests that this person lacks the technical ability to handle the security of his/her internet-connected devices, the respondent still mentioned a 3 on a scale of 0 to 5. This was also observed during the interviewees with respondents 8 and 9. For example, respondent 9 explains that: *'Before I received the email I already made the plan to contact someone who has more knowledge about this [computer security] as we already have issues that need to be resolved and I can't do it'*. Nevertheless, this interviewee responded with a 3 as well on a scale of 0 to 5.

Fourth, the type of end-user is identified. The interviewees were asked if they use their internet connection for business or private purposes. In TG1 & TG2 only one of the 12 interviewees explained that the internet connection was used for business purposes instead of private use. More interesting is TG3, the group which is assumed to consist mainly of firewall owners. Of the three participants interviewed, two affirmed that they utilised their internet connection for commercial purposes. The remaining three contacted customers, who did not participate in an interview, mentioned the company's name when answering the phone, implying that their internet connection is also utilised for business operations. 45 KPN customers from TG3 were not approached for an interview. However, based on the available contact information, 23 out of 45 of these customers could be classified as probable businesses as they provided email addresses associated with business entities, indicating that these IP addresses are being employed by a small-scale enterprise. Due to the limited sample size of three interviewed participants from TG3, it is not possible to arrive at a definitive conclusion regarding the composition of the entire target group. Nonetheless, this observation strongly suggests that this target group largely consists of businesses that use these IP addresses.

Last, interviewees were asked if they had experienced security issues before in regard to their smart devices or their internet connection in general. Of the 15 respondents, only one (respondent 10) reported a prior incident, while the remainder confirmed that they had not experienced any prior security issues, and this was their first encounter with such a matter raised by the ISP.

6.1.2. Vulnerable device identification

The first step for a vulnerability notification to work is to identify the devices which cause the problem. The notification explained that there are different types of devices that could be used for TCP reflective amplification attacks. The email states that it could either be a device part of the internet infrastructure (e.g. firewall or router/modem) or a smart device (e.g. energy monitor, alarm systems, camera, building management system). Due to the wide variety of devices which can be abused, it was expected that not all end-users were not able to identify the vulnerable device.

Number of smart devices in households

In general, end-users have multiple smart devices in their homes. On average the interviewees own 4 smart devices when looking at all interviewed end-users (all target groups combined). It is important to note here that this is based on the number of devices named by the interviewees. Not all interviewees do have a clear view of what smart devices are. For instance, respondent 2 noticed: *'My husband does*

not think it [BMS] has been updated, I myself had never realised that it was using WiFi'. To obtain an accurate number, interviewees were given prompts and suggestions during the interviews to identify most of the smart devices. However, it is still likely that the number of smart devices in the household of the interviewees is larger. The number of smart devices named by the interviewees is illustrated in figure 17.

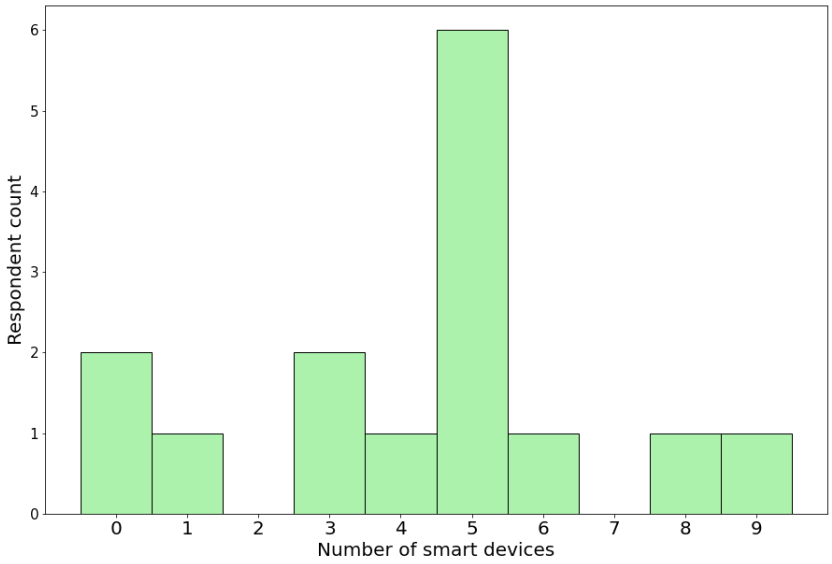

Figure 17: Number of smart devices in the homes of interviewees

Difficulties in vulnerable device identification

Interviewees show that they have difficulties in identifying the right device based on the vulnerability notification. Only 4 out of 11 interviewees were able to locate the device based on the email. However, during the interview, two of these interviewees mentioned that they looked at the device which is vulnerable but also mentioned numerous other devices which they think could cause problems. For example, respondent 5 said: *'After reading the email, I did not know what the device could be, you would be surprised how much of that junk you have in your house these days. But the first devices at which a took a look were the energy monitor and the camera system as those were explicitly mentioned.* Respondents 8 and 15 were the only interviewees who were able to directly pinpoint the vulnerable device. However, both respondents explain that their reasoning for pinpointing the right device was not based on the information provided in the notification. Respondent 8: *'We do not have smart devices in our home except our washing machine. This device must be the cause of the problem as we opened all ports to be able to use it, I know it is not safe but it is the only way to make that thing work'*. Despite the fact a washing machine was not explicitly mentioned in the email, this respondent was able to identify the device based on the fact there is only one smart device in the household. Respondent 15 was able to identify the vulnerable device by arguing that it must be the new device which was recently installed shortly before he/she received the email notification. An explanation for this is the fact that respondent 15 is first seen in the data by Shadowserver during the research period, while most interviewees from TG1 are reported on before the start of the research period.

The other seven respondents were not able to identify the device which causes the problem. Respondents frequently report attempting to identify the vulnerable device based on their own reasoning rather than relying on the examples of devices provided in the notification. Out of these interviewees, 4 thought of a device which was not mentioned in the email. Interestingly, all these 4 interviewees had the same line of reasoning for identifying the device. These end-users explained that they recently bought a new device and shortly after they received the notification, thus concluding that the new device should be the one causing problems. 2 of these 7 interviewees simply did not take a look at any of their devices as they did not have the knowledge about computer security (respondent 9) or did not think the problem to be urgent (respondent 7). The last interviewee (respondent 3), wrongly assumed that the

router was the vulnerable device. As the router was the first type of device that was mentioned in the email, it could be that this interviewee did not look at the other possible devices mentioned because of not reading the complete text.

The device that respondents tend to pinpoint after the notification is bought by themselves instead of provided by a third party, such as a router provided by an ISP. 10 of the 11 respondents explained that they bought the devices themselves. However, two respondents mentioned that they are not responsible for managing and maintenance of the device. These were respondents 2 and 4, which owned a vulnerable building management system and alarm system respectively. While they bought these systems by themselves they were not responsible for installing and managing the system. The only respondent which did not buy the vulnerable device was respondent 7 who happens to be the only interviewee from TG1 who uses the internet connection for business purposes. This respondent explained that: *'I was not here yet when this heating system was installed or bought so I don't know'*.

6.1.3. Performed actions

Respondents in TG1 were asked about their performed actions after reading the notification. The different actions are visualised in table 5. It could be that an interviewee performed multiple actions, therefore the total number of executed actions is higher than the number of interviewees.

Table 5: Performed actions after receiving notification

Performed action	Respondent count
Check software update (smart) device	5
No action performed	3
Ask someone for help	2
Set password (smart) device	1
Execute port scan	1
Install anti-virus software	1

The most mentioned action taken by the interviewees was in line with the vulnerability notification stating to look for updates for the suggested (vulnerable) devices. 5 of the 11 interviewees checked a device for updates which they assumed to be causing the problem. However, none of these 5 interviewees did install a software update as the device showed that the device was already up-to-date. For example, respondent 1 stated: *'after reading the email I visited the website of the manufacturer of the energy monitor to see if they have something mentioned about this issue there, but there was nothing there and it explained that my device is already up-to-date'*. The action of checking for software updates was the intended behaviour of the interviewees to do after receiving a notification. Nonetheless, only respondents 8 and 15 checked the device which is perceived as vulnerable to TCP reflective amplification attacks as found in chapter 5. The other three respondents checked for updates on devices which are likely not the problem, thus not fully addressing the issue discussed in this thesis.

Comparatively, other interviewees (n=3) mentioned that they did not perform any action in response to the notification. The reasoning differs among these three interviewees for not doing anything. Respondent 4 explained that: *'The information provided was really generic, also when clicking on the exclamation mark in the email you get directed to a standard KPN commercial page. This made that whole email lose all its power as it made me think that nothing serious is happening on my internet connection'*. Respondent 7 did not perform any actions as he/she did not receive the email because it was sent to another email address. The person who received the email explained that he/she deleted the email simply because: *'I'm only responsible for the financial part of the internet connection and not the technical part'*. The one responsible for the internet connection (respondent 7) was never notified by this person. Lastly, respondent 11 mentioned that he/she did not perform any actions because he/she thought that sending a large file was the *'trigger'* for the ISP to send an email. Additionally, there were two interviewees that asked someone for help after receiving the email. Even though these two interviewees performed no further steps, they are not classified as 'No action performed' as potentially they would perform the desired behaviour after getting a consult from someone else. As for this thesis, interviewees were called two days after receiving the email, these participants were possibly not given enough time to get help. Therefore, their action is labelled as 'Ask someone for help'.

Finally, it should be noted that some interviewed participants (n=3) took actions aimed at enhancing the security of their internet connection but did not address the issue of TCP reflective amplification attacks. For instance, Respondent 5 secured his/her energy monitor by setting a password. The respondent stated, *'I decided to use the password functionality for the energy monitor. I already knew that it was possible to use a password but I didn't think it was scary for that device to be open to the internet.* This action did not fix the broken TCP protocol, thus the device can still be used for launching TCP reflective amplification attacks. Nonetheless, the amplification factor is substantially smaller when the password functionality is used, as explained in 5. Respondent 2 and Respondent 6 implemented measures to enhance the security of their network. Specifically, Respondent 2 installed a virus scanner, while respondent 6 performed a port scan to identify vulnerabilities in his/her network. Although these actions represent steps in the right direction towards securing their own internet connection, they did not address the issue of a device with a broken TCP protocol in their network. As such, the desired impact was not realised. However, it is noteworthy that these measures had the positive side effect of enhancing the overall security of their internet connection.

6.1.4. Smart device update behaviour

All the respondents (TG1, TG2 and TG3) were asked about their security behaviour in regard to updating their (smart) devices. Updating the software or hardware of devices with broken TCP protocols or middleboxes is a solution for end-users to prevent their device from being used for launching TCP reflective amplification attacks.

The respondents were surveyed regarding their approaches to updating their (smart) devices. Given that the interviewees often possess numerous devices within their household, different devices may require distinct updating protocols. For instance, automatic updating may be suitable for one device, while manual checking may be necessary for another. As such, interviewees may describe multiple practices for updating their (smart) devices. These different practices are visualised below in table 6.

Table 6: Different practices for interviewees to update their (smart) devices

	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	
Update automatically	X		X	X	X	X			X	X	X	X	X	X	X	12
Update after notification	X	X	X					X	X	X		X		X	X	9
Update after manual checking						X				X						2
Not updating devices							X									1
Device managed by third party		X		X									X			3

Interviewees (n=12) rely on automatic updates for their smart devices or think that their smart devices are automatically updated. It is hard to say if people know this really is the case. Even if end-users know that probably not all devices are updated automatically, they are not aware of the risks or tend to take the risk. For example, respondent 11 states: *'I always set the devices of the big brands to be done automatically and I have to be honest then I also blindly assume that this is also done automatically. I have to say that sometimes I find out by chance that that doesn't happen, but I don't actively check it'.* Respondent 4 explains that: *'I also have devices that I do not update, I wonder if my router amplifier can be automatically updated. I assume all equipment is just automatically updated. I do not check whether updates are available for the smart equipment but assume that it just happens'.* Only respondents 4 and 11 exclusively stated that they rely completely on automatic updates. While the other interviewees, who also mentioned that updates occur automatically, also explained that they do not depend solely on this feature.

Another update practice for smart devices, as mentioned by interviewees (n=9), is to install updates in response to notifications indicating the availability of an update. Only respondent 8 indicated that he/she exclusively relies on these notifications. This can be explained by the fact that the same respondent stated that he/she only has one smart device within their household. 7 out of 9 interviewees indicated during the interview that they not only rely on updating their smart devices after receiving notifications but also possess devices that update automatically. This observation might suggest a heightened level of awareness regarding which devices can be updated in their households, as they can differentiate

between varying updating practices. In contrast, those interviewees who rely solely on one updating practice may have a comparatively lower level of awareness.

Moreover, interviewees (n=2) do check actively if their smart devices are up to date with the latest software or firmware. However, this is probably often on an ad-hoc basis. For example, respondent 6 stated: *'I don't look at it every day, but for example when I have time during the weekend, I go through everything. That way I keep it all a bit up to date.* Another interviewee, respondent R10, mentioned: *'Coincidentally, I did update the energy monitor a few weeks ago because new firmware was available'.* but later the same respondent explained that: *'I was already busy with it [energy monitor] and then I saw somewhere that a new update was available'*.

According to interviewees (n=3), the task of updating some smart devices falls to a third party, even if end-users own the devices themselves. These devices were found to be part of alarm systems and building management systems that were installed by a third party. Therefore, these interviewees suggested that the responsibility for updating these devices rests with the third party.

One interviewee explained that he/she does not install updates for their smart devices. Respondent 7 argued: *'I don't know if it happens automatically, I don't think so. I actually think the functionality is not given. I myself never actually look at the updates for the smart devices hanging in the church, I have to admit very honestly'.* This lack of updating can be explained by the fact that this interviewee does not use the internet connection for private purposes and thus possibly feels a reduced amount of responsibility.

Incentives for updating

The interviewees expressed multiple incentives for updating smart devices in their households. These different incentives are displayed below in table 7, including the number of times the interviewees named these incentives for updating their smart devices.

Table 7: Reasoning interviewees for updating smart devices

	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	
Externally prompted updating	X	X	X		X			X	X	X	X	X	X		X	11
Security	X	X				X	X		X		X			X		7
Happens automatically				X					X			X			X	4
New functionalities		X														1
Device does not work							X									1

Interviewees (n=11) emphasised that their decision to update their smart devices was influenced by external incentives. However, they did not cite a specific rationale for doing so but rather mentioned updating in response to receiving a notification or based on the fact that updates are always suggested or recommended. For example, respondent 10 stated: *'Actually, from all devices I have at home I get a ping when new firmware is available, in fact, the software and firmware of every device mentioned is up to date because I already do that'.* Similarly, respondent 6 explained: *'... and then I usually get a message that an update is available. And I do that by default if devices support it'.* Often interviewees also mention that they update their smart devices for security reasons, this is mentioned by 7 interviewees. Respondent 9 remarked this by stating: *'Because it [updating] is recommended. Lately, you are getting scared of all those hacks that are being performed so I try to secure everything as well as possible.* Another interviewee, respondent 1, who also mentioned security as a reason, used different reasoning for his/her explanation: *'I update that because it's safe and because it's recommended. I'm pretty into it [topic of computer security] myself so I'm pretty sharp on that.'*

Other interviewees (n=3) explained that (some) smart devices are updated automatically and did not further specify their reasoning for updating. For example, respondent 4 states: *'... I just assume everything happens automatically'*. Only one interviewee mentioned that he/she updates for new functionalities, however, respondent 2 was talking about his/her tablet when stating: *'I always do the updates myself when I get a message stating that it [tablet] would work better and that it is [tablet] better secured.* Finally, it is worth noting that only one interviewee provided an explanation for not updating their smart

devices, stating that he/she does not update his/her devices as long as the devices continue to perform their intended functions satisfactorily: *'... to be honest it [smart devices] works and I didn't really think about it [not updating] being a security issue so for the church that was the most important thing and for me that was the reason I didn't consciously do this [updating] in the past.'* However, the respondent expressed his/her concern about the security issue during the interview and mentioned that he/she will take action for security reasons afterwards. Therefore this is also included as a type of reasoning for this respondent in table 7.

Incentives not for updating

Besides expressing the incentives for updating their smart devices, interviewees were also asked for their reasoning for not updating their smart devices. The different responses are visualised below in table 8.

Table 8: Reasoning interviewees for not updating smart devices

	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	
Unaware of update capability		X	X	X	X	X			X		X	X				8
No reason for not updating	X							X		X			X	X	X	6
Interrupts daily business					X						X					2
It already works							X									1

The interviewees (n=14) explained that they do their updates unless they don't know a certain device can be updated. Nonentheless, respondent 7, argued that as long the smart devices performed satisfactorily he/she did not see the urge for updating the device. Additionally, it was contended by respondents (n=2) during the interviews that end-user can refrain from updating their smart devices due to disrupting their device's performance. Nevertheless, both interviewees clarified that they only delay the updating process rather than entirely abstain from it. Interestingly, respondent 3 mentioned during the interview: *'If I'm going to update and reset then I'm afraid I'll end up in a pandemic of adjustments where I'll need my son again so I'm very hesitant to do that'.* As this interviewee did not explicitly mention this when asked about his/her reasoning for not updating, it is not included in table 8. Nonetheless, this perspective offers valuable insights into why the participant may choose to postpone or overlook the installation of a specific update.

Lastly, interviewees were asked whether they think a notification by the ISP helps them with updating their smart devices. The respondents (n=9) explained that notification might help them with reminding them to update. Respondent 5 explains: *'I think so. It's a warning anyway and you don't want there to be a problem so that's good'.* On the other hand, the other interviewees (n=6) do not perceive any benefit from receiving a notification. The majority of these individuals clarified that they already make an effort to keep their smart device software up to date and thus make such notification by the ISP redundant. For example, respondent 4 argues: *No I already do it [updating], that is why this would appear redundant to me'.* One interviewee explained that he/she was not sure whether it would help him/her or not.

6.1.5. End-user notification perception and feedback

At the end of the interview, respondents were asked about their thoughts in regard to the service of the ISP to reach out to its customers with vulnerable devices. The respondents (n=14) clarified that they are rather positive about the ISP notifying them about vulnerabilities. Nevertheless, a single interviewee, identified as respondent 4, expressed his/her dissatisfaction with this service stating: *'... if you send such an email again, I will immediately throw it away because it has far too little information'.* Despite the overall positive consensus among the interviewees regarding the service provided by the ISP to notify its customers about vulnerabilities, the perception expressed by respondent 4 about the inadequate and limited amount of information is frequently cited during the interviews. 6 out of 11 respondents were unsure of the notification's effectiveness, believing that it lacked essential information needed for remediating the vulnerability. Table 9 provides a visual representation of the number of times a shortcoming in the notification is mentioned by an interviewee, including the specific pieces of information that interviewees perceived as missing. This table is based on the interviews with respondents from

TG1 as the other groups did not receive a notification.

Table 9: Missing information in notification according to interviewees

Type of missing information	Respondent count
What device causes the issue	5
IP address	4
Date of issue detection	3
No missing information	3
Urgency of the issue	2
Port number	2
Not sure if information is missing	1
MAC address	1
Type of abuse	1

The types of information reported as missing can be categorised into two distinct groups: those related to the vulnerable device and those concerning the specific nature of the vulnerability.

Interviewees mostly ask for more information related to the vulnerable device itself. This is shown by the number of times ”IP address” and ”What device causes the issue” as shown in table 9. This is not surprising as in the notification itself various potentially vulnerable devices were mentioned. However, the IP address is explicitly shown on top of the email that the interviewees received. They either have missed it or would like to have info on the local IP address of the device causing issues. As respondent 5 stated in the interview: *'... There [top of the email] the IP address is indicated in green and it took me a while before I could click on it too. For the rest, nothing was indicated about what kind of device it was, only an IP address. Later I could click on it and see what kind of device is behind it. I didn't see anything in the mail about which device it was'.* Other interviewees did also mention that they would like to have more information on the vulnerable device such as the port number (respondents 6 and 10) and the MAC address (respondent 5).

The interviewees also expressed a desire for more information regarding the vulnerability itself. Specifically, three of them stated that they would like to know the timing of the vulnerability, such as when the ISP first detected it and how long it has been an issue. This suggests that the participants are not only interested in being informed about the existence of the vulnerability, but also about its timeline and duration. By providing this type of information, individuals may be better informed to identify which devices are vulnerable. During the interviews, it became apparent that some participants attempted to link the timing of the notification email to a newly acquired device or a change made to a particular device. Therefore, offering more specific details about the timing and duration of the vulnerability could enable users to identify the vulnerable device. Furthermore, 2 interviewees would like to have more information on the urgency of the problem as this would help them with potentially mitigating it. For example, *'I would like to know how serious this problem is because that system [alarm system] is under the management of another company for maintenance. Then I must also be able to tell them more about the problem to solve the issue'.*

Other respondents (n=5) explained that they think there is no additional information needed for performing the actions they deem necessary for mitigating the vulnerability or that they are not sure if they need additional information.

In addition to the perceived lack of information, another noteworthy observation when examining end-users perceptions of vulnerability notifications is the topic of trust. Out of the 15 interviews conducted, only three respondents discussed the topic of trust, citing their need to verify the authenticity of the email or phone call from the ISP before proceeding. One respondent, respondent 8, expressed distrust and sought to verify the call by contacting the ISP's customer service before proceeding with the interview. Similarly, respondents 2 and 3 also initially expressed their lack of trust and emphasised the importance of verifying the notification before taking any actions.

It is important to note that the observation regarding trust may be biased, as individuals who trust the notification, email, or phone call from the ISP may be more inclined to participate in the interviews. This

bias is illustrated by the responses of other customers who were contacted but declined to participate in the interviews, often citing their lack of trust in either the call or the email they received. These responses were not included in this thesis as they were not obtained through interviews.

6.1.6. Observed differences among target groups

Differences and similarities were noted among the three target groups regarding end-user characteristics and update behaviour. Due to the small sample size of interview participants, the present findings cannot be used to draw generalisable conclusions regarding the similarities and differences observed. Nonetheless, the following two discussion points offer valuable insights into understanding TCP reflective amplification attacks in an ISP network.

A difference can be seen when looking at the characteristics of end-users in TG1, TG2 and TG3. As discussed in sub-chapter 6.1.1, in TG3 more than half of the end-users use them for business purposes instead of personal use. This could explain why firewalls (middleboxes), were detected in TG3 but not in TG1 and TG2, where none of the devices were classified as middleboxes. It makes sense for businesses to operate a larger and more complex internet network compared to consumer end-users which do not need more advanced IT infrastructure such as firewalls as these are already built into most routers, modems and PC's.

Interestingly, the sole interviewee from TG2 revealed during the interview that he/she had not performed any updates or made any other changes to the vulnerable device in question, which happens to be the same device that is commonly found in TG1. Despite this interviewee's lack of action, his/her IP address no longer appears in abuse feeds, indicating that the previously broken TCP protocol has been resolved. It is unclear why this respondent no longer appears in the abuse feeds, given that the device in question is still accessible over the internet and no patch or update has been released by the vendor, as discussed in sub-chapter 6.2.

6.2. Manufacturer

In this thesis, one manufacturer is contacted who is responsible for the energy monitor that is frequently regarded as a vulnerable device, due to a broken TCP protocol in TG1 and TG2 as outlined in chapter 5. The main takeaways from this interview are presented in this sub-chapter.

First, the manufacturer was unaware of the problem that was caused by the device or the fact that the device has a broken TCP implementation. However, the manufacturer was known of the problem (vulnerable middleboxes and TCP reflected amplification attacks) but did not expect that the device was able to cause the amplification as found in this thesis. The manufacturer argued: *'I think I've read about it. I know vaguely but to be honest I don't really know exactly what it means. It has something to do with an exploit where you can somehow generate a DDoS attack, but I don't know exactly how it works'.* Later during the interview the manufacturer mentions when talking about receiving the notification email: *'... I thought that we would neatly complete the handshake in terms of IP stack, so I thought that the specific problem in the mail was not relevant to us so it [email notification] must be something of a routine notification that a port is open'.*

The second takeaway is the size of the company and the number of devices it produces. As it was already assumed based on the website of the manufacturer that the size of the manufacturer would be relatively small. This is confirmed during the interview where is stated: *'...it's actually a bit of a niche device that we make'.* Additionally, during the end-user interview, respondent 11 said the following about the energy monitor: *'because it [energy monitor] is quite an open source thing made by some amateurs, so that is usually not so safe I think.*

Thirdly, the manufacturer employs a single method for detecting security issues, namely by exposing the device to the internet and monitoring for potential threats that could pose a risk to the device or its users: *'... it is a kind of honeypot. To see what types of attacks are being released to see if we can learn something from them. We look at what is attracted by the open device and does that cause a problem or not. By looking at a log file we can see what is being released and whether it [the energy monitor] continues to work properly.* Next to this, the manufacturer argues that he/she *'reads some things once in a while that could be relevant for us or the device'.*

The fourth takeaway is that the manufacturer only communicates new firmware updates via a message on a blog on the website: *'We have a blog on our website where we announce that there is a new update where we also mention what things have been solved or improved. The firmware updates themselves can be done on the app.* Nevertheless, this method may not reach every customer. This is explained by respondent 6 who mentioned the following during the interview: *'...I don't think that [energy monitor] is something for which new software is released every year [...] I also don't know how I would know that there is new software because I never get an email from those people'.*

This statement by the interviewee suggests that there is an information asymmetry between the manufacturer and the end-users of the energy monitor. The manufacturer itself expressed that he/she was surprised by the fact that end-users have the device open to the internet and that they often do not set a password. The manufacturer explains that: *'... it is possible, people who do have knowledge of it can open the web interface [...] in principle, this is not something we recommend doing'.* Besides when talking about the option for setting a password, the manufacturer said: *'The manual states that a password is recommended [...] it actually surprises me that there are so many [end-users without password] while they do have it open on the internet, which is not really a responsible thing to do'.*

6.3. Applying the COM-B model

The COM-B model is used to analyse the aforementioned findings of end-user behaviour in regard to the vulnerability notification and their smart device update behaviour. The themes discussed are linked to the theoretical constructs as presented in chapter 3.3. The study's findings are analysed and discussed based on the three components: capability, opportunity, and motivation. Furthermore, a detailed examination is conducted of the sub-components that make up these three primary components. This analysis forms the basis for answering the main research question in terms of end-user behaviour in chapter 8.

6.3.1. Capability

The capability component refers to all internal factors that may impact an individual's ability to perform the necessary steps outlined in a vulnerability notification. In general, the capability of the interviewees to perform the desired actions after receiving the notification depends on the person itself and therefore varies among the interviewees. As was shown in chapter 6.1.2 and 6.1.3. People tend to perform various different actions after receiving the vulnerability notification, often neglecting the action as stated in the notification itself: checking for updates for a vulnerable smart device. It was found that overall the capability component can be divided into two different aspects that affect the end-user behaviour, the capability to identify the vulnerable device and the capability to perform the intended actions as stated in the notification.

Results show that the first step to identifying the vulnerable device may be difficult as end-users can own multiple smart devices which could be the potential source of the problem. Only two interviewees were able to identify the vulnerable device and perform the desired action of checking for a software or firmware update for that device.

However, when looking at the interviewees' (smart) device update behaviour, end-users do seem to have the capability to install software or firmware on their (smart) devices. Based on the results in chapter 6.1.4, it was found that the participants reported keeping their smart devices up to date, with only one individual reporting otherwise. These findings suggest that there are no apparent barriers related to the capability component preventing end-users from updating their devices, indicating that they are both physically and psychologically capable of performing updates as needed.

Physical capability As discussed in chapter 4.4.1, the physical capability is not directly asked during the interview as it is known that the vulnerable device is physically present at the home of the interviewee and that it is not likely that an interviewee has no computer or smartphone to make changes to the vulnerable device in terms of software or firmware updates. Nevertheless, none of the respondents indicated that they lacked some sort of physical capability to perform the actions after reading the vulnerability notification.

Psychological capability Even though this is not directly queried during the interview it was found that in general, the interviewees show that they possess the psychological capability to perform the

intended actions. The instructions provided in the notification merely involved two actions: identifying the vulnerable device and verifying the presence of software or hardware updates. The end-users did not encounter any technical difficulties as the instructions did not necessitate complex reasoning. This can also be seen in the high average level of perceived skills in computer security and the fact that there are multiple respondents with a background in ICT, suggesting that the interviewees have the psychological capability to perform the intended behaviour. Moreover, this observation was substantiated by the interviewees, where none of the participants, except one, reported any difficulties in comprehending the steps outlined in the notification. Respondent 2 explicitly stated: *'... I just find it [vulnerability notification] difficult to interpret because I have too little knowledge in that area'.*

Nonetheless, some of the results suggest that respondents lack the psychological capacity to perform the intended actions. As discussed in chapter 6.1.1, although the average level of perceived computer security skills is high among the participants, the interviews revealed that some respondents mistakenly overestimate their competence in this area. For instance, respondent 9 explicitly indicated that the email was clear. Nevertheless, the same interviewee mentioned that assistance from another individual was still required to carry out the recommended steps. Moreover, the observed reasoning behind the respondents' actions may indicate a lack of psychological capacity, given that their responses frequently deviated from the recommended course of action outlined in the vulnerability notification. It is difficult to conclusively attribute this solely to a lack of psychological capacity. It could also be due to not fully reading the notification or disregarding the advice provided, under the assumption that they know better than the recommendations outlined in the vulnerability notification. For example, as explained in chapter 6.1.2, interviewees often attempted to connect the timing of the received notification to a recently purchased or installed device, even if the device was not explicitly mentioned in the email.

6.3.2. Opportunity

The opportunity component can be described as all the external factors that affect individuals performing the steps described in the vulnerability notification. It appears that end-users faced a lack of opportunity to address the vulnerability due to insufficient information provided in the vulnerability notification and the absence of updates. Specifically, the notification did not provide enough details for users to identify which device was affected, and in addition, no update was available to remedy the vulnerability. As a result, end-users were unable to take the necessary steps to address the vulnerability in their devices.

Physical opportunity During the interviewees it became apparent that the interviewees had the physical opportunity to perform a behaviour as all the interviewees, except respondent 7, had received and read the vulnerability notification. Respondent 7, did not receive the email because his/her email was not known by the ISP, and the receiver did not forward the mail to this respondent.

However, besides the information provided in the notification, and on the ISP website, multiple respondents needed additional information and tools to perform the steps. As explained, in chapter 6.1.5, interviewees often asked for more information on the vulnerability and the vulnerable device during the interviews. This suggests that the end-users lacked the psychical opportunity in the form of information provided by the ISP. Furthermore, some respondents expressed that they needed tools (e.g. a port scanner) for performing their actions, as is described in 6.1.3. In addition, the interview with the manufacturer in chapter 6.2 revealed that there was no update available to address the identified vulnerability in the energy monitors. As a result, the interviewees lacked the physical opportunity to perform an update on these vulnerable devices.

An additional noteworthy finding is that some interviewees either stated that they had not taken any action yet due to perceiving a lack of urgency, or expressed during the interview that they needed more information regarding the severity of the vulnerability. This may indicate that these individuals were either unable or unwilling to make time to address the issue, thereby lacking the necessary physical opportunity in terms of time to perform the steps outlined in the notification. It is important to note that this may not solely be attributed to physical opportunity, as the decision to make time for taking action is also influenced by the motivation component.

Social opportunity Overall, the interviewees did not mention any social norms or expectations that may have influenced their behaviour when carrying out the (recommended) actions. The only thing

that can be observed in the findings relating to this component is that respondents (n=2) asked for help to perform the actions, thus addressing the social opportunity component. Respondent 2 further noted that they checked the authenticity of the email with their son before taking any action, to ensure that it was genuinely from the ISP and could be trusted. This suggests that the behaviour of the interviewee may have been influenced by the cyber security norms and perception of the son.

Another important note can be made regarding the social opportunity component in terms of distrust by end-users towards the vulnerability notification. Although the interviewees did not display any social norms or expectations that could have impacted their behaviour, it was noted that customers who were contacted but did not wish to participate frequently demonstrated a lack of trust towards the interviewer. This distrust may affect the behaviour as distrust likely resulted in not performing any steps after receiving a vulnerability. As this distrust is fed by the news, people's opinions and other forms of information on cyber security, it is important to provide a detailed explanation of the problem, a plausible reason for notifying the end-uses, and, if applicable, offer a solution, to help build trust (Hennig et al., 2022).

6.3.3. Motivation

The motivation component can be described as all the processes happening in the brain of the end-users when receiving a vulnerability notification and performing certain remediating steps. Generally, it was observed that end-users are motivated to update their smart devices when they are aware that updates are available. However, when faced with a vulnerability notification, end-users may not always be able to take the necessary steps to address the issue, indicating that their motivation to act may be impacted by other factors such as the lack of information provided.

Reflective processes A large share of the interviewees, as explained in 6.1.5, indicated that they believed the steps described in the notification were not useful for them as they lacked the necessary information for remediating the vulnerability. Moreover, the study frequently observed participants seeking additional information during the interviews. This behaviour may be attributed to the fact that individuals with questions or concerns about the vulnerability or notification were more prone to participate in the research. Nevertheless, this suggests that end-users generally behave rationally. They analyse the available information and attempt to perform the correct behaviour in response to the vulnerability notification at hand.

Another interesting observation, as detailed previously, is the tendency for individuals to associate the timing of a vulnerability notification with a recently purchased or installed device when attempting to identify the vulnerable device. This often leads to the incorrect device being identified as vulnerable. However, this method of reasoning may be effective in identifying the correct vulnerable device in the future, as the devices of the KPN customers contacted for this research were already reported by abuse data at the beginning of Shadowserver's reporting, indicating a significant vulnerability period. Therefore, this reasoning is not useful for correctly identifying the vulnerable device in these cases. Nevertheless, as exemplified in chapter 6.1.2, Respondent 15 demonstrates how this reasoning approach can be utilised to identify the vulnerable device when the individual has recently purchased or installed it.

Additionally, this study discovered that a significant portion of the vulnerable devices in TG1 consists of energy monitors that are frequently accessible via the internet without requiring a password. According to the manufacturer's explanation provided in 6.2, this practice is not standard. This indicates that end-users may not fully comprehend the potential risks associated with having an internet-accessible device without a password. Although respondents appear to comprehend the necessary actions to take after receiving a notification, this finding implies that they may not fully grasp the risks associated with internet-accessible devices. Therefore, they may not fully understand the potential consequences of failing to follow the advice provided by an ISP or other entity aimed at enhancing computer security.

This observation of interviewees not understanding the risks was further observed in chapter 6.1.3, as there were 3 respondents who did not perform any action. This suggests that do not understand, or are aware of, the risks of not performing the steps. An additional element that could have influenced the conduct of end-users in failing to take action or taking the incorrect action is their reasoning that a

particular device is managed by a third party, and this is not their responsibility to update, as discussed in chapter 6.1.4. This thought process can hinder the end-users in recognising the vulnerable device and undertaking the appropriate steps in response to a vulnerability notification.

Automatic processes Overall, it was observed that all interviewees had a positive perception of receiving a vulnerability notification, and considered it a valuable service provided by the ISP. However, as described previously, shortcomings in the notification used for this study were identified by the interviewees, which may affect their general attitude towards vulnerability notifications. For instance, respondent 4 expressed dissatisfaction with the service during the interview.

Moreover, various interviewees explained that they do not think that the notification is effective as they already keep all their devices up to date, thus making the notification redundant. This is contradicting the finding of interviewees stating that they perceive the service of vulnerability notifications in general as positive, suggesting that it helps them with keeping their devices up to date.

Other emotions that were seen during the interviews were surprise and curiosity. Respondent 10 explained that he/she was surprised by the email as he/she already keeps all his devices secure and up to date. The curiosity was shown by the many questions that respondents have about the vulnerability itself and the vulnerable device, as is shown in 6.1.5. It is important to note that this finding may be biased, as it is possible that end-users who are more interested in learning about the vulnerable device or vulnerability may be more willing to participate in the interviews.

7

Discussion

This thesis has explored the issue of vulnerable middleboxes and IoT devices which can be abused for TCP reflective amplification attacks in terms of the type of vulnerable devices and end-users (updating) behaviour in the network of an ISP. In this chapter, findings are put into context to describe what they mean and represent. Moreover, this chapter discusses the limitations of the research which have an effect on the validity of the results.

The results indicate that the problem of vulnerable devices which can be abused for TCP reflective amplification attacks are actually two different problems which work in a similar way and can be used for the same type of DDoS attack by generating amplification over TCP protocols. Bock et al. (2021) already explained that a small share of the vulnerable devices are not middleboxes causing the amplification but does not further specify the characteristics. This research found that these are normal consumer IoT devices with broken TCP protocols which can be exploited. However, a large emphasis can still be laid on middleboxes. When looking at the absolute numbers, the amount of vulnerable consumer IoT devices in the KPN network is very limited. It was found that in the summer of 2022 security patches for firewalls were released, remediating around 70 per cent of the total population of vulnerable devices. After this, what remains is the limited number of around 15 detected vulnerable IoT devices per week. Based on this one can say that the large problem, these vulnerable firewalls, is already fixed, and only some broken IoT devices are left which are coincidentally identified when scanning for vulnerable middleboxes. However, this research also found substantial amplification factors created by these IoT devices, exceeding amplification factors from commonly used UDP-based protocols (Bjerre et al., 2022). This shows that even though in absolute numbers the amount of vulnerable devices is very limited, the devices that are vulnerable can be used for generating large amplification and thus do form a problem for any potential victim.

These two different problems are differently located on the vulnerability life cycle as shown in chapter 2.4. This is visualised below in figure 18. The problem of vulnerable middleboxes is located somewhere on the right side of the figure. The findings from this research suggest that patches for vulnerable middleboxes, specifically firewalls, in the KPN consumer network are released and installed, remediating the vulnerability. However, it can not be concluded that for all middleboxes vulnerable to TCP reflective amplification attacks patches are released. Conversely, the vulnerability within IoT devices is situated somewhere on the left side of the vulnerability life cycle. This thesis discovered the problem in IoT devices but most certainly no patches are released yet, as is shown by the manufacturer interview in chapter 6.2.

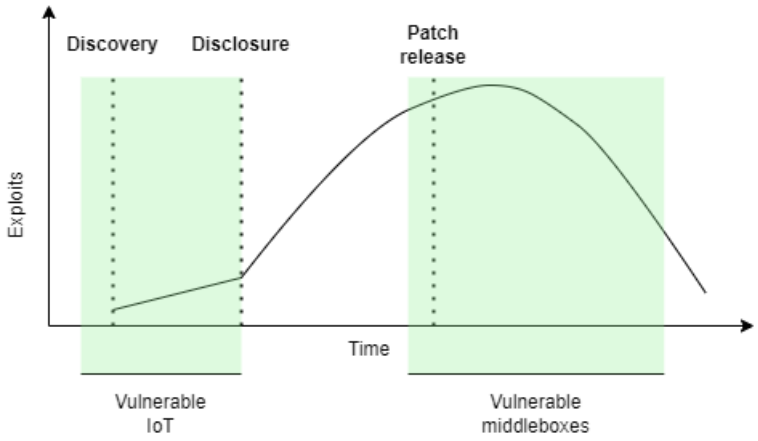

Figure 18: Problem of middleboxes and IoT vulnerable to TCP reflective amplification attacks displayed on the vulnerability life cycle

Chapter 4.3 shows that while the consumer market of KPN has a relatively small number of devices, the overall number of vulnerable devices across the entire KPN network, including the business and wholesale markets, is significantly higher. Additionally, the decline in vulnerable devices following firewall security patches in the summer of 2022 is not as apparent in the business and wholesale markets compared to the consumer market as shown in chapter 5. Although the business and wholesale markets were not in the scope of this thesis, this observation, combined with the finding that consumer IoT devices are vulnerable to TCP reflective amplification attacks, suggests that there is a large number of vulnerable IoT devices in the KPN network. As attackers do not differentiate between markets, this emphasises the presence of the issue in the KPN network, and targeting the remediation of devices outside of the consumer market may have a greater impact on remedying the issue as a whole within the KPN network.

This study found that the opportunity component of the COM-B model plays an important role in the behaviour of end-users during the process of updating smart devices, as it enables end-users to take the necessary steps after being notified of a vulnerability. End-users are both motivated and capable of keeping their (smart) devices up to date but needed comprehensive information to pinpoint the correct device and to understand the need for remediation. This observation is in line with other researchers who studied end-user security behaviour using the COM-B model (Rodriguez et al., 2022; van der Kleij et al., 2021).

Furthermore, the findings from this research indicate that end-users often would like to have extensive information on the vulnerability when receiving a vulnerability notification. It is worth noting that the observation made in this study is based on a limited sample size, which restricts the generalisability of these findings to a larger population. Nonetheless, these are in line with research that suggests that vulnerability notification should be comprehensive in order to be effective (Çetin et al., 2016; Durumeric et al., 2014). However, it can be seen as contradictory to studies by other researchers such as Forget et al. (2016), who state that vulnerability notifications should be simple and easy to interpret. An explanation for the identified need for additional information by end-users may be explained by the high average level of perceived skills in computer security by the interviewees. It can be expected that people with more knowledge about this topic are more interested in extensive (technical) information on the vulnerability. However, this research also identified end-users with a lesser understanding of computer security who did not express the desire for more information. This suggests that a vulnerability notification should be tailored fit to end-users in order to enable them to identify the vulnerable device and perform the intended steps.

The interviews with end-users also brought to light that certain vulnerable devices are under the management of a third party, which highlights a security issue discussed in chapter 2.1. It can be unclear who bears responsibility for securing these at-risk devices (Maple, 2017). If a third party manages the devices, end-users may not feel accountable for updating them. Therefore, while vulnerability notifications are a critical component of addressing security concerns by ISP's, they are not always sufficient

to prompt end-users to update their devices. The aforementioned findings highlight a factor that could hinder effective vulnerability notification, not only for notifying consumers about vulnerable devices susceptible to TCP reflective amplification but also for other types of security incidents or abuse, such as malware infections.

This study has shown that within the consumer network of KPN, there are vulnerable devices which can not be remediated by updating the device as there is no update available that fixes the issue. As a result, it is worth considering whether vulnerability notifications via email can be a helpful tool for informing and assisting end-users with their devices. Another notification method, such as walled garden notifications, may be more effective to remediate the vulnerability by simply forcing the end-user to remove the vulnerable device from the internet. Previous studies found that these types of notifications are more effective for malware remediation than email notifications, suggesting that it may be an interesting option for ISP's to adopt for remediating vulnerable consumer IoT devices which can be abused for TCP reflective amplification attacks (Cetin, Ganan, Altena, Kasama, et al., 2019). While updating the devices is a suggested solution that remediated middleboxes, it may not be practical for end-users with vulnerable consumer IoT devices. Ultimately, it is the responsibility of the device manufacturer to ensure that their products perform a complete TCP handshake before transmitting data over the internet and to provide updates to ensure that this handshake works correctly at all times. Motivating end-users to update their smart devices may not necessarily resolve the vulnerability if the manufacturer is either unaware of the issue or is (unintentionally) neglecting to develop updates to address it. In such cases, the vulnerability will remain present.

Limitations of research

It is important to highlight the different limitations of the research and used methods. This will provide perspective on the validity of the results which are presented. Some limitations of the method have already been discussed in chapter 4.6 and will not be further elaborated on in this chapter.

The first limitation of this study is the generalisability of the research findings based on the interviews. This can be attributed to two different factors. First, as discussed in chapter 4.6 the sample size of interviews was restricted due to the small population size of KPN consumer market. In particular, when examining TG2 and TG3, a limited number of interviews were conducted, which restricts the ability of this research to extrapolate its findings to the entire population of KPN. For TG1, all customers (n=18) were approached, of whom 11 participated in the study, suggesting that generalisability is less of a concern for this particular target group. Nevertheless, it is vital to consider this limitation when interpreting the results, as the quantitative outcomes cannot be applied to the entire KPN population or to a broader context. The second factor is the fact that this research only analyses the issue within the consumer market of a single Dutch ISP. Although the qualitative findings of the interviews can be extended to the population of KPN consumers with vulnerable IP addresses that can be exploited for TCP reflective amplification, it is challenging to extrapolate them to a wider population such as other ISP's in the Netherlands or in the world.

Another significant limitation of this research relates to the use of semi-structured interviews as the research method. While chapter 4.6 discusses response and social desirability biases, there are additional limitations worth mentioning. The findings presented in this thesis were derived from the interpretation of information provided by the interviewees. Although the semi-structured format of the interviews allowed for more in-depth questioning in cases of ambiguity, it is reasonable to assume that the researcher may have introduced biases during the interviews and the subsequent interpretation of results, in addition to any biases exhibited by the interviewees themselves. Moreover, the results from the interviews were analysed by using the COM-B model. However, it is worth noting that not all components of the model were directly queried during the interviews such as the social opportunity component. Therefore the researcher's interpretation of the results of the interviews could have influenced the findings from this analysis. Another risk of semi-structured interviews is that every interview progress slightly differs from the prepared interview protocols. In some cases, interviewees are more invested in the interview and decide to share their actions, thoughts and perception in greater detail compared to other interviewees or vice versa.

The third limitation is that this research solely relies on the identified IP addresses based on Shadowserver abuse feeds for vulnerable middleboxes. This research shows that there are vulnerable IoT

devices which respond to the same scanning method as Shadowserver uses for vulnerable middlebox identification. However one can not be sure that this does identify all the vulnerable IoT devices with broken TCP protocols. Therefore, it may be that this research does not give a complete overview of the vulnerable devices in the consumer market of KPN and thus does not fully picture the issue. It could be that there are more devices in the KPN network with broken TCP protocols which can be abused for TCP reflective amplification attacks. Furthermore, this research assumes that devices are remediated when they are not showing up anymore in the abuse feeds by Shadowserver. However, is hard to be completely sure that the device is remediated

The last limitation that can be discussed in regard to characterising vulnerable devices. As discussed before, by manually scanning vulnerable IP addresses, this research was only able to verify the characteristics of the vulnerable devices which were vulnerable at the moment of scanning. Because this research tries to characterise the devices over a longer period of time, it is unable to be fully sure about the identified characteristics of devices which were vulnerable prior to the start of this research. Shadowserver started reporting on the vulnerability in April 2022 while this research started in November 2022. It could be that IP addresses were last seen somewhere in for example June but this research tries to characterise them a few months later in December. In this meantime, people could have changed the internet setup within their homes potentially leading to characterising the incorrect device.

8

Conclusion

In this chapter, the answers to the main research question and sub-questions are presented. First, the conclusion with the answers to the research questions is introduced (§8.1). Following, the recommendations for KPN and other ISP's are presented (§8.2). Third, the societal and scientific relevance of the findings and conclusions of this research are discussed (§8.3 and §8.4). Last, possible directions for future work are proposed (§8.5).

8.1. Answer to research questions

This research aimed to explore the issue of vulnerable middleboxes and IoT devices which can be abused for TCP reflective amplification in terms of the type of vulnerable devices and end-users (updating) behaviour in the consumer network of an ISP. A mixed-method approach is chosen, focusing on a single case, for exploring this issue. Network scans are performed analysing the vulnerable IP addresses within the network of consumer customers of KPN to identify the vulnerable devices. Besides this, fifteen end-users, and one manufacturer of these devices are interviewed, to collect data on their updating behaviour and characteristics. This approach is used in order to formulate an answer to the main research question:

To what extent can we characterise the issue of TCP reflective amplification in the network of an ISP?

The main research question is divided into four sub-questions. In the following sections, the conclusions to the sub-questions are discussed based on the findings which are shown in chapter 5 and 6.

SQ1: What does the population of vulnerable devices look like in the consumer network of an ISP?

The first sub-question aims to characterise the vulnerable device in the network of KPN. In total 79 unique consumer IP addresses have been identified in the network of KPN with vulnerable devices which can be abused for TCP reflective amplification attacks between the 25th of April 2022 and February 19th 2023. It was found that part of this group is remediated and part of the group is not remediated during the research period between January 16th and February 19th 2023. The group of vulnerable devices (n=18), which was not remediated, mainly consisted of one type of energy monitor (n=11) and other types of consumer IoT devices. The other found types are, building management systems (n=3), alarm systems (n=1), washing machines (n=1), WiFi extenders (n=1) and digital video recorders (n=1). These devices do not fit the description of middleboxes which is remarkable given the fact, that the identified IP addresses are based on internet scans looking for vulnerable middleboxes. This research found that (consumer) IoT devices also respond to these scans and can be used for launching TCP-reflected amplification attacks. However, compared to vulnerable middleboxes, IoT devices respond with large URL landing pages to the internet scans because of a broken TCP protocol which does not complete the three-way TCP handshake before sending the page. This is different from vulnerable middleboxes which respond to forbidden URL requests with large block pages because of the policy rules within that middlebox. Nevertheless, both types of devices can be used for creating the same

effect, creating significant amplification factors which can be used for launching DDoS attacks aimed at potential victims. While the result is the same, the working of the vulnerability is different between IoT devices and middleboxes.

The group of devices which were considered remediated during the research period (n=61) consisted of a variety of different types of devices. However, the difference with the group as described above is that both consumer IoT devices and middleboxes, in the form of firewalls, were identified. In total 16 of these devices were characterised as firewalls. This research was unable to identify the vulnerable devices for about half of the IP addresses in this group (n=31). However, the findings from this research suggest that these are mainly firewalls but this can not be concluded. Additionally, the consumer IoT devices that were identified in this group consist of the same type of energy monitor (n=5), building management systems (n=2), routers (n=3), digital video recorders (n=2) and (web-)servers (n=2).

SQ2: What are the characteristics of vulnerable device end-users?

The second sub-question strives to define the characteristics of end-users in an ISP's network in terms of demographics and technical expertise. In total 31 KPN customers were contacted, of which 15 participated in interviews which partially covered end-users characteristics. During the interviews, a relatively high average age of end-users was observed at 60 years old, compared to the average age of the KPN customer population being 55. Furthermore, it was found that often male customers are responsible for the middleboxes or IoT devices in the homes where the vulnerable devices are located. The majority of interviewees (n=12) identified themselves as male, while the remaining interviewees (n=3) identified as female. However, the gender distribution derived from the interviews only reflects the gender of the participants interviewed. Therefore, it is not appropriate to extrapolate this distribution to all end-users of the vulnerable devices. It is possible that other people in the household, who were not interviewed, could also use the vulnerable device. Nonetheless, the study did not explicitly inquire about this during the interviews. On the other hand, the gender distribution mentioned above suggests that males are primarily accountable for managing the security issues of vulnerable devices, regardless of whether or not they are the only users of the device in the household.

In addition, the study found that, on average, the interviewees had a high perception of the level of their computer security skills. Despite perceiving themselves as having a high level of computer security skills, some respondents displayed a lower level of such skills during the interview. This suggests that end-users do sometimes overestimate their ability to adequately secure their (smart) devices. On the other hand, other respondents mentioned that they have a background in ICT and show a high level of computer security skills during the interview. Therefore, it can be concluded that, in terms of the level of computer security skills, a wide variety of end-users operate a vulnerable device which can be abused for TCP reflective amplification attacks. A crucial point to note is that the observation of lower levels of computer security skills is based on the researchers' interpretation of the interviews and has not been tested. Therefore, no definite conclusion can be drawn on this topic.

A more notable trend emerged when examining whether end-users utilise the internet connection for business or personal purposes. Even though this research focused on the consumer market within KPN, it was found that multiple customers use consumer internet subscriptions for their businesses. As can be expected, findings from this research suggest that IP addresses with vulnerable consumer IoT devices are often used for private purposes, while IP addresses with vulnerable firewalls (middlebox) are often used for business purposes. All of the IP addresses which were found to be used for business purposes are considered to be remediated during the research period (January 16th till February 19th 2023) and thus are not vulnerable anymore to be abused for TCP reflective amplification attacks. Only one device was considered not remediated during the research period and was used for business purposes.

SQ3: What does the updating behaviour of vulnerable device end-users look like in the consumer network of an ISP?

The third sub-question aims to capture the (smart) device updating behaviour of end-users. In general, from the findings, it can be concluded that if people know a device can be updated they say that they perform the update, however, there are likely devices which escaped the attention of the end-users. A large share (n=12) of the interviewees indicated that the updates of their (smart) devices happen

automatically or that they think that these updates happen automatically. It is difficult to confirm whether this is indeed true, as seven out of the 15 interviewees mentioned that they were unaware of certain devices and whether they could be updated. Moreover, interviewees (n=9) stated that they update their devices after they receive a notification reminding them of the availability of an update. Only one of the 15 interviewees stated that he/she did not update the devices connected to the internet connection.

End-users show a willingness to perform updates if they know there is an update for a device available. The respondents (n=10) indicated that they performed updates either because they receive a notification or because it was ”recommended”. These findings indicate that vulnerable device end-users are typically prompted by external factors. When asked for the reasoning for not updating the device, all interviewees except one (n=14) explained that there are no reasons for not performing updates unless they are not aware of a device that can be updated to a new software or firmware version. Only one interviewee showed that he/she did not perform updates as the device already worked accordingly.

SQ4: What is the effect of a vulnerability notification on vulnerable device end-users in terms of performed actions and perception?

The fourth, and last sub-question, tries to identify the effect of a vulnerability notification on end-users of the vulnerable device which can be abused for TCP-reflected amplification attacks. In general, it can be seen that end-users do not perform the intended steps of vulnerability notifications which can be attributed to two different factors. Firstly, it is difficult for end-users to identify the correct device that is vulnerable based solely on the information provided in the notification. With the average number of (smart) devices in a household being four among the interviewees, there are multiple potential devices that could be at risk. Secondly, end-users feel that the vulnerability notification did not provide sufficient information, leading them to take actions that deviate from the intended course or no action at all. End-users perceived the vulnerability notification send for this research as vague or generic and thus would like to have more information on the vulnerability itself and the vulnerable device in question.

Most of the interviewees (n=14) perceived the service of the ISP informing them of a vulnerability as positive. Nevertheless, respondents (n=9) indicated that the email helped them with keeping their (smart) devices up to date with the latest software and hardware version. On the other hand, interviewees (n=6) do not think a vulnerability notification is effective in helping them with updating their smart devices. The main reason for this is that end-users already make an effort to keep their smart device up to date, thus such notification has no effect on them.

Main conclusions

To conclude, the answer to the main research question can be formulated based on the answers to the sub-questions. The issue of vulnerable devices that can be exploited for TCP reflective amplification attacks can be classified into two distinct issues. First, the problem of vulnerable middleboxes with certain security policies which makes them over-responsive to forbidden URL’s and can be used for this type of attack. These middleboxes are often used for business purposes. Second, the issue of consumer IoT devices with broken TCP protocols which respond to requests from the internet without completing the TCP handshake. These devices are often used by consumers. However, the problem of vulnerable middleboxes seems to be solved in the network of KPN, as all the vulnerable firewalls that have been identified are considered remediated. This is in line with expectations as vendors of firewall hardware released patches remediating the vulnerability in the summer of 2022. Thus, as of now, only the problem of consumer IoT devices with broken TCP protocols remains.

This key finding of this thesis is particularly noteworthy. While previous studies by researchers, such as Bock et al. (2021), Q. Li et al. (2023), and Nosyk et al. (2022), have focused on middleboxes when investigating TCP reflective amplification, this study has uncovered that consumer IoT devices, and potentially other internet-connected devices with broken TCP protocols, can be exploited in a similar manner as middleboxes for launching DDoS attacks. Similar to the vulnerability in middleboxes, as shown by Bock et al. (2021), consumer IoT devices can generate significant amplification factors that exceed those of commonly known UDP-based amplifiers, highlighting the severity of the vulnerability in these devices.

While updates may have resolved the issue with vulnerable firewalls, updating consumer IoT devices does not necessarily address the vulnerability present in the devices that have been identified. In fact, it is quite likely that manufacturers are unaware of the broken TCP protocol in some of the devices they produce, and thus there are no updates available to remedy the situation. This thesis reveals that for KPN, a significant proportion of the vulnerable consumer IoT devices are energy monitors, and there is at the moment of writing no update available to address the vulnerability in these devices. This is because the manufacturer was unaware of the issue and thus could not have provided a remedy.

Despite the absence of updates addressing the issue of broken TCP protocols in consumer IoT, end-users are willing to install updates and vulnerability notifications can be a tool for an ISP to use to inform end-users of the vulnerability and achieve a certain behaviour by these end-users. The research results indicate that vulnerability notifications can be useful for end-users as they serve as reminders to update their devices. The majority of end-users tend to perform updates when they are aware that updates are available and often do so in response to a notification. However, the study results suggest that a vulnerability notification should include a comprehensive description of the vulnerable device and the specific vulnerability, enabling end-users to accurately identify the device in question and perform the necessary software or firmware update.

The findings from the interviews were analysed based on the COM-B model. Overall, this analysis revealed that end-users are both motivated and capable of keeping their (smart) devices up to date. However, they often lack the physical opportunity to do so due to the generic nature of the information provided in vulnerability notifications. Additionally, often there is no update available as the device is already up to date, amplifying the issue of physical opportunity. This conclusion reinforces the importance of the opportunity component in successfully performing the intended steps outlined in the notification. Moreover, this analysis highlights that the interviewees frequently possess both the physical and psychological capability to execute the steps outlined in a vulnerability notification. Furthermore, regarding the motivation component, the interviewees showed that they mainly use reflective processes in driving their behaviour, rather than automatic processes, suggesting that they act rationally when acting to the vulnerability notification.

8.2. Recommendations

From the results and conclusions of this study, recommendations can be formulated for KPN and other ISP's. Due to the exploratory nature of this research, the findings can not be used for specific policy recommendations. Instead, the findings provide background information to understand the problem of interest and what ISP's or other involved actors may use to make better-informed (policy) decisions concerning vulnerable middleboxes and/or vulnerable IoT devices which can be abused for TCP reflective amplification attacks. At the moment of writing, the research findings indicate that relying solely on email notifications to encourage end-users to update their devices may not be sufficient to address the vulnerability. Hence, it may be worth exploring alternative notification methods such as walled garden notifications, which could be a more effective approach for ISP's to prompt end-users to remove the vulnerable device or disconnect it from the internet, thus remediating the vulnerability.

Although findings suggest vulnerability notifications may not feasible for remediating the vulnerability, they have the potential to be a successful strategy in the future. This study's results suggest that individuals possess both the ability and the incentive to inspect their vulnerable devices for updates upon receiving a notification from their ISP, under the condition that it contains comprehensive information. Notifying end-users about the vulnerability of their consumer IoT devices with malfunctioning TCP protocols may serve as an effective approach to address the vulnerability by prompting them to update their devices. However, at the moment of writing, another notifying approach such as walled garden notification may be more effective. To further clarify this suggestion for ISP's, the following recommendations may be provided.

First, ISP's could focus the content of their vulnerability notification on the findings of this research. Even though Shadowserver scans search for vulnerable middleboxes, currently only vulnerable consumer IoT devices are identified by these scans. This suggests that information communicated towards consumer customers should be focused on these devices instead of middleboxes. Furthermore, the results from this research show that customers show on average a high level of skill in the field of

computer security and would like to have extensive information. Therefore, it can be recommended to provide additional (technical) information on a website page or attachment, to inform the customers who would like to have more information. This additional information could provide information on the vulnerable device for example in the form of a list of known devices including brand name or type. Furthermore, additional information on the vulnerability and the impact of the vulnerability may increase the perceived urgency by end-users to act.

Second, focus on the manufacturer of vulnerable devices in the notification. This research shows that manufacturers are not always aware of the vulnerability and therefore patching does not seem to be a one size fits all solution. To avoid the burden of an ISP contacting multiple manufacturers of vulnerable devices, it could be suggested to encourage end-users to notify the manufacturer about the issue of broken TCP protocols. When end-users take the initiative to contact the responsible manufacturer to remediate the vulnerability, the manufacturer will be automatically notified about the vulnerability in their devices. However, it requires end-users to provide some information to the manufacturer which may not be practical for end-users who do not understand anything about the message that the ISP is trying to convey.

Third, start notifying business customers. This was out of the scope of this research but nonetheless, it was found that a large proportion of the vulnerable devices in the network of KPN were in the business market. An attacker does not see the differences between business and consumer customers as the devices can have a similar effect by generating large amplification factors. When an ISP would like to start solving the issue itself within its network, this is the place where the most can be gained.

8.3. Societal contribution

The findings and conclusions of this research provide background information on the issue of devices vulnerable to TCP reflective amplification. This guides ISP's, like KPN, on how to notify end-users of the vulnerable devices to remediate the problem. However, the impact of the findings can also be discussed in a broader sense.

First, this research studies end-user updating behaviour and presents several insights on this behaviour in terms of practices used by end-users and reasons for (not) updating. Performing updates can solve other security-related issues apart from the issue of interest in this thesis. Therefore a better understanding of people's behaviour in relation to updating their smart devices can be seen as valuable from a broader perspective. The national Dutch government has been conducting awareness campaigns to encourage Dutch residents to install updates and ensure the security of their smart devices (Rijksoverheid, 2021). Nevertheless, it was found that these campaigns do not have the desired effect (Cammaert et al., 2022). The findings of this study contribute to this topic by suggesting that, in addition to raising awareness, external factors such as notifications can also play a vital role in motivating end-users to install software or firmware updates. The results of this thesis indicate that people are likely to perform updates in response to a notification from the manufacturer or vendor that an update is available. Therefore, improving communication from the manufacturer about new updates could be an effective approach to encourage end-users to perform updates, instead of just raising general awareness about their importance.

Furthermore, as discussed in chapter 1.5, this research is relevant to society due to the potential impact of DDoS attacks using TCP reflective amplification on victims. However, the results of this study indicate that the number of vulnerable devices in the KPN consumer network is limited, meaning that remedying these will not significantly reduce the potential for DDoS attacks exploiting this vulnerability. Conversely, the number of vulnerable devices in the business market of KPN is much greater, highlighting the importance of addressing the issue in that market. Nevertheless, the characteristics of vulnerable devices discovered in this study can still be valuable in remediating devices in the business market, stressing the relevance of the findings of this research.

8.4. Scientific contribution

In addition to discussing the societal contribution, the significance of this study's findings and conclusions can also be examined in terms of their scientific contribution. This research makes contributions to three distinct fields, which are elaborated on below.

First of all the research provides novel findings in regard to TCP reflective amplification attacks using middleboxes or other devices. First of all, this research characterised the devices that can be abused for TCP reflective amplification attacks, which is not done before (Bock et al., 2021; Nosyk et al., 2022). Moreover, this research identifies a new technique of launching TCP reflective amplification attacks involving IoT devices, which is similar to the attack technique described by (Bock et al., 2021) but different from the focus on SYN-ACK floods emphasised in previous research on TCP reflective amplification such as Kührer et al. (2014a), Kührer et al. (2014b), and Mohana Priya et al. (2014). According to the findings of this study, TCP reflective amplification attacks can be initiated by exploiting vulnerable IoT devices through URL requests, which result in the transmission of the internet landing page of that device.

Second, this research contributes to the field of vulnerability notifications by ISP's. While this study does not focus on the effectiveness of such messages it does provide background information on the effect of vulnerability notifications on end-users to mitigate devices vulnerable to being abused for TCP reflective amplification attacks. This study can be seen as the starting point for future studies on this topic.

Last, this study addresses and contributes to the topic of end-user (updating) behaviour. The subject of end-user updating behaviour has been studied in the past (Sarabi et al., 2017). Nevertheless, the findings and conclusions presented in this thesis contribute to this field of knowledge. This study examined end-users behaviour related to updating their smart devices and responding to vulnerability notifications, identifying underlying themes and reasoning. This thesis focuses on end-user behaviour in terms of vulnerability mitigation, specifically the updating of their smart devices to prevent vulnerabilities, such as broken TCP protocols, from being abused. This is a departure from previous studies, such as Rodriguez et al. (2022) and Rodríguez et al. (2021), which primarily investigated end-user behaviour in terms of malware cleanup. At the moment of writing, this thesis is the first to analyse end-user behaviour in regard to TCP reflective amplification attacks and vulnerability notifications.

8.5. Future work

Based on the before-mentioned limitations of this study, possible directions for future research can be described. This study was performed by only looking at consumer customers of a dutch ISP. This resulted in a limited number of conducted interviews. Future studies can be performed among a larger group of participants to increase the generalisability of the results. This could be done for example by looking at different ISP's in the Netherlands, Europe or worldwide. This would especially increase the generalisability in regard to the results about the end-user security behaviour in terms of updating behaviour and performed actions after receiving vulnerability notifications as these are research topics which are interesting on their own.

Studying a larger population size also enables future research to validate the findings in this study through statistical analysis. By the means of a survey, a quantitative perspective can be laid on the qualitative findings as shown in this research. This would allow researchers to statistically analyse the correlations between different factors which may or may not affect the behaviour, perception and actions of end-users after receiving a vulnerability notification for the issue of interest. This thesis could provide researchers with a basis for the survey setup as different factors are identified which affect the end-users.

Another possible direction for future research is to study the issue of vulnerable middleboxes and IoT devices in the business market. Notifying business customers likely requires another approach than notifying consumer customers in terms of content. However, this was out of the scope of this study and thus needs further research to analyse the differences between consumer and business end-users in terms of behaviour and perception. The results from the end-users interviews already suggest that the behaviour and perception of business end-user differ from consumer end-users as was shown by the businesses that were in the population of interest of this research.

Research could focus on the vulnerability notifications and how different channels or content may affect the end-users in terms of updating behaviour, perception and actions. While the goal of this thesis was to explore the issue, the next step would be to study whether such notification and in what form, would be effective in remediating the vulnerability. By performing similar research in the future, it may be able

to test the effectiveness of notifying end-users of vulnerable devices as at that time there are updates available which should able end-users to remediate the vulnerability. Additionally, this study only sends one vulnerability notification to each contacted end-user. It may be that multiple notifications may affect the behaviour and perception of end-users, something which is not covered within this thesis.

Lastly, this study discovered that certain devices are not vulnerable anymore but no update had been released to fix the issue suggesting that something else remediated the vulnerability. It is unclear how this natural remediation has occurred especially as the manufacturer, in the case of identifying the energy monitors, has not released an update fixing the issue at the moment of writing. Further studies could try to provide answers in regard to this unclarity as that would provide insights into a potential remediation solution.

Bibliography

ACM. (2023). ACM Telecommonitor derde kwartaal 2022. https://www.acm.nl/nl/publicaties/acm-telecommonitor-derde-kwartaal-2022

Agrawal, N., Zhu, F., & Carpenter, S. (2020). Do You See the Warning? *Proceedings of the 2020 ACM Southeast Conference*, 260–263. https://doi.org/10.1145/3374135.3385314

Al Alsadi, A. A., Sameshima, K., Bleier, J., Yoshioka, K., Lindorfer, M., van Eeten, M., & Gañán, C. H. (2022). No Spring Chicken: Quantifying the Lifespan of Exploits in IoT Malware Using Static and Dynamic Analysis. *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, 309–321. https://doi.org/10.1145/3488932.3517408

Alani, M. M. (2014). OSI Model. In *Springerbriefs in computer science.* (pp. 5–17). Springer, Cham. https://doi.org/10.1007/978-3-319-05152-9{_}2

Altena, L. (2018). *Exploring Effective Notification Mechanisms For Infected IoT Devices* (Doctoral dissertation). Delft University of Technology.

Anand, P., Singh, Y., Selwal, A., Alazab, M., Tanwar, S., & Kumar, N. (2020). IoT Vulnerability Assessment for Sustainable Computing: Threats, Current Solutions, and Open Challenges. *IEEE Access*, *8*, 168825–168853. https://doi.org/10.1109/ACCESS.2020.3022842

Anell, S., Grober, L., & Krombholz, K. (2020). End User and Expert Perceptions of Threats and Potential Countermeasures. *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 230–239. https://doi.org/10.1109/EuroSPW51379.2020.00038

Arbaugh, W., McHugh, J., & Fithen, W. (2000). Windows of vulnerability: a case study analysis. *Computer*, *33*(12), 52–59. https://doi.org/10.1109/2.889093

Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, *54*(15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010

Beel, J., & Gipp, B. (2009). Google Scholar's ranking algorithm: an introductory overview. *Proceedings of the 12th international conference on scientometrics and informetrics (ISSI'09)*, 230–241.

Bjerre, A., Westh, A. P., Villefrance, E., Haque, A. S. M. F. A., Andersen, J. B., Helgogaard, L. K., & Anagnostopoulos, M. (2022). A Wide Network Scanning for Discovery of UDP-Based Reflectors in the Nordic Countries. In *Nordsec 2022: Secure it systems* (pp. 176–193). https://doi.org/10.1007/978-3-031-22295-5{_}10

Bock, K., Alaraj, A., Fax, Y., Hurley, K., Wustrow, E., & Levin, D. (2021). Weaponizing Middleboxes for TCP Reflected Amplification. *30th USENIX Security Symposium (USENIX Security 21)*, 3345–3361. https://www.usenix.org/conference/usenixsecurity21/presentation/bock

Boehmer, J., LaRose, R., Rifon, N., Alhabash, S., & Cotten, S. (2015). Determinants of online safety behaviour: towards an intervention strategy for college students. *Behaviour & Information Technology*, *34*(10), 1022–1035. https://doi.org/10.1080/0144929X.2015.1028448

Bougie, R., & Uma Sekaran. (2019). *Research methods for business: A skill building approach*. John Wiley & Sons.

Bouwmeester, B., Rodríguez, E. R. T., Gañán, C. H., van Eeten, M., & Parkin, S. (2021). "The Thing Doesn't Have a Name": Learning from Emergent Real-World Interventions in Smart Home Security. *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 493–512.

Braun, V., & Clarke, V. (2021). One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology*, *18*(3), 328–352. https://doi.org/10.1080/14780887.2020.1769238

Cammaert, M., & Aarts, J. (2022). *Campagne-effectonderzoek DOE JE UPDATES WINTER 2021/2022* (tech. rep.). DVJ Insights. https://open.overheid.nl/Details/ronl-f783e95bacff59bdd111de57e32cef7a479c28ff/1?hit=4&thema=c_ee06665e&informatiesoort=c_3300f29a&organisatie=mnre1045&page=1&count=10#panel-gegevens

Carpenter, B., & Brim, S. (2002). *Middleboxes: Taxonomy and Issues* (tech. rep.). https://doi.org/10.17487/rfc3234

Cetin, O., Ganan, C., Altena, L., Kasama, T., Inoue, D., Tamiya, K., Tie, Y., Yoshioka, K., & van Eeten, M. (2019). Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. *Proceedings 2019 Network and Distributed System Security Symposium*. https://doi.org/10.14722/ndss.2019.23438

Cetin, O., Ganan, C., Altena, L., Tajalizadehkhoob, S., & Van Eeten, M. (2019). Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Networks. *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, 326–339. https://doi.org/10.1109/EuroSP.2019.00032

Cetin, O., Ganan, C., Korczynski, M., & van Eeten, M. (2017). Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. *16th Workshop on the Economics of Information Security (WEIS 2017)*. http://www.infosecon.net/workshop/downloads/2017/pdf/Make_Notifications_Great_Again:_Learning_How_to_Notify_in_the_Age_of_Large-Scale_Vulnerability_Scanning.pdf

Çetin, O., Hanif Jhaveri, M., Gañán, C., van Eeten, M., & Moore, T. (2016). Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity*, *2*(1), 83–98. https://doi.org/10.1093/cybsec/tyw005

Chen, F., Huo, Y., Zhu, J., & Fan, D. (2020). A Review on the Study on MQTT Security Challenge. *2020 IEEE International Conference on Smart Cloud (SmartCloud)*, 128–133. https://doi.org/10.1109/SmartCloud49737.2020.00032

Chen, J. (2020). Risk communication in cyberspace: a brief review of the information-processing and mental models approaches. *Current Opinion in Psychology*, *36*, 135–140. https://doi.org/10.1016/j.copsyc.2020.06.006

CIRT Akamai. (2022). TCP Middlebox Reflection: Coming to a DDoS Near You. https://www.akamai.com/blog/security/tcp-middlebox-reflection

Cloudflare. (2022). What is a DDoS attack? https://www.cloudflare.com/learning/ddos/what-is-a-ddos-attack/

Conzola, V. C., & Wogalter, M. S. (2001). A Communication–Human Information Processing (C–HIP) approach to warning effectiveness in the workplace. *Journal of Risk Research*, *4*(4), 309–322. https://doi.org/10.1080/13669870110062712

Creswell, J. W. (2014). *Research Design: Qualitative, Quantitative and Mixed Methods Approaches* (4th ed.). SAGE Publications, Inc.

Das, R., Karabade, A., & Tuna, G. (2015). Common network attack types and defense mechanisms. *2015 23nd Signal Processing and Communications Applications Conference (SIU)*, 2658–2661. https://doi.org/10.1109/SIU.2015.7130435

De Carli, L., & Mignano, A. (2021). Network Security for Home IoT Devices Must Involve the User: A Position Paper. In *Fps 2020: Foundations and practice of security* (pp. 20–28). https://doi.org/10.1007/978-3-030-70881-8{_}2

Douligeris, C., & Mitrokotsa, A. (2004). DDoS attacks and defense mechanisms: classification and state-of-the-art. *Computer Networks*, *44*(5), 643–666. https://doi.org/10.1016/j.comnet.2003.10.003

Durumeric, Z., Li, F., Kasten, J., Amann, J., Beekman, J., Payer, M., Weaver, N., Adrian, D., Paxson, V., Bailey, M., & Halderman, J. A. (2014). The Matter of Heartbleed. *Proceedings of the 2014 Conference on Internet Measurement Conference*, 475–488. https://doi.org/10.1145/2663716.2663755

ENISA. (2019). Cybersecurity Culture Guidelines: Behavioural Aspects of Cybersecurity. https://www.enisa.europa.eu/publications/cybersecurity-culture-guidelines-behavioural-aspects-of-cybersecurity

European Commission. (2017). Resilience, Deterrence and Defence: Building strong cybersecurity for the EU. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52017JC0450&from=EN

Fagan, M., Khan, M. M. H., & Buck, R. (2015). A study of users' experiences and beliefs about software update messages. *Computers in Human Behavior*, *51*, 504–519. https://doi.org/10.1016/j.chb.2015.04.075

Forget, A., Pearman, S., Thomas, J., Acquisti, A., Christin, N., Faith, L., Egelman, S., Harbach, M., & Telang, R. (2016). Do or Do Not, There Is No Try: User Engagement May Not Improve Security Outcomes. *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 97–111.

Fortinet. (2022). FortiOS - TCP Middlebox Reflection. https://www.fortiguard.com/psirt/FG-IR-22-073

Frei, S., May, M., Fiedler, U., & Plattner, B. (2006). Large-scale vulnerability analysis. *Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense - LSAD '06*, 131–138. https://doi.org/10.1145/1162666.1162671

Fruchter, N. H. (2019). *Enhancing ISP-consumer security notifications* (Doctoral dissertation). Massachusetts Institute of Technology.

Gordon Lyon. (2008). *Nmap Network Scanning*. Insecure.Com, LLC.

Guest, G., Bunce, A., & Johnson, L. (2006). How Many Interviews Are Enough? *Field Methods*, *18*(1), 59–82. https://doi.org/10.1177/1525822X05279903

Gusmeroli, S., Piccione, S., & Rotondi, D. (2013). A capability-based security approach to manage access control in the Internet of Things. *Mathematical and Computer Modelling*, *58*(5-6), 1189–1205. https://doi.org/10.1016/j.mcm.2013.02.006

Haney, J. M., & Furman, S. M. (2021). Work in Progress: Towards Usable Updates for Smart Home Devices. In *Stast 2020: Socio-technical aspects in security and trust* (pp. 107–117). https://doi.org/10.1007/978-3-030-79318-0{_}6

Heer, T., Garcia-Morchon, O., Hummen, R., Keoh, S. L., Kumar, S. S., & Wehrle, K. (2011). Security Challenges in the IP-based Internet of Things. *Wireless Personal Communications*, *61*(3), 527–542. https://doi.org/10.1007/s11277-011-0385-5

Hennig, A., Neusser, F., Pawelek, A. A., Herrmann, D., & Mayer, P. (2022). Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications. *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–8. https://doi.org/10.1145/3491101.3519847

Hoque, N., Bhattacharyya, D. K., & Kalita, J. K. (2015). Botnet in DDoS Attacks: Trends and Challenges. *IEEE Communications Surveys & Tutorials*, *17*(4), 2242–2270. https://doi.org/10.1109/COMST.2015.2457491

Horton, J., Macve, R., & Struyven, G. (2004). Qualitative research: experience in using semi-structured interviews. In *The real life guide to accounting research: A behind-the-scenes view of using qualitative research methods* (pp. 339–358). Elsevier Science.

Islam, U., Muhammad, A., Mansoor, R., Hossain, M. S., Ahmad, I., Eldin, E. T., Khan, J. A., Rehman, A. U., & Shafiq, M. (2022). Detection of Distributed Denial of Service (DDoS) Attacks in IOT Based Monitoring System of Banking Sector Using Machine Learning Models. *Sustainability*, *14*(14), 8374. https://doi.org/10.3390/su14148374

Jalali, S., & Wohlin, C. (2012). Systematic literature studies: database searches vs. backward snowballing. *Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement - ESEM '12*, 29–38. https://doi.org/10.1145/2372251.2372257

Kallio, H., Pietilä, A.-M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, *72*(12), 2954–2965. https://doi.org/10.1111/jan.13031

Kar, S., Chakravorty, B., Sinha, S., & Gupta, M. P. (2018). Analysis of Stakeholders Within IoT Ecosystem. In *Advances in theory and practice of emerging markets* (pp. 251–276). https://doi.org/10.1007/978-3-319-78378-9{_}15

Katwala, A. (2019). Your old router is an absolute goldmine for troublesome hackers. https://www.wired.co.uk/article/router-wifi-security-settings

Kaur Chahal, J., Bhandari, A., & Behal, S. (2019). Distributed Denial of Service Attacks: A Threat or Challenge. *New Review of Information Networking*, *24*(1), 31–103. https://doi.org/10.1080/13614576.2019.1611468

Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and Other Botnets. *Computer*, *50*(7), 80–84. https://doi.org/10.1109/MC.2017.201

Kührer, M., Hupperich, T., Rossow, C., & Holz, T. (2014a). Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks. *8th USENIX Workshop on Offensive Technologies (WOOT 14)*.

Kührer, M., Hupperich, T., Rossow, C., & Holz, T. (2014b). Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. *23rd USENIX Security Symposium (USENIX Security 14)*, 111–125. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhrer

Li, F., Bailey, M., Durumeric, Z., Czyz, J., Karami, M., Mccoy, D., Savage, S., & Paxson, V. (2016). You've Got Vulnerability: Exploring Effective Vulnerability Notifications. *25th USENIX Security Symposium (USENIX Security 16)*, 1033–1050. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li

Li, F., Ho, G., Kuan, E., Niu, Y., Ballard, L., Thomas, K., Bursztein, E., & Paxson, V. (2016). Remedying Web Hijacking. *Proceedings of the 25th International Conference on World Wide Web*, 1009–1019. https://doi.org/10.1145/2872427.2883039

Li, Q., Huang, H., Li, R., Lv, J., Yuan, Z., Ma, L., Han, Y., & Jiang, Y. (2023). A Comprehensive Survey on Ddos Defense Systems: New Trends and Challenges. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4363418

Li, Y., & Liu, Q. (2021). A comprehensive review study of cyber-attacks and cyber security; Emerging trends and recent developments. *Energy Reports*, *7*, 8176–8186. https://doi.org/10.1016/j.egyr.2021.08.126

Lin, J., Yu, W., Zhang, N., Yang, X., Zhang, H., & Zhao, W. (2017). A Survey on Internet of Things: Architecture, Enabling Technologies, Security and Privacy, and Applications. *IEEE Internet of Things Journal*, *4*(5), 1125–1142. https://doi.org/10.1109/JIOT.2017.2683200

Livingood, J., Mody, N., & Reirdan, M. (2012). Recommendations for the Remediation of Bots in ISP Networks. https://www.rfc-editor.org/rfc/rfc6561

Lyu, M., Sherratt, D., Sivanathan, A., Gharakheili, H. H., Radford, A., & Sivaraman, V. (2017). Quantifying the reflective DDoS attack capability of household IoT devices. *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 46–51. https://doi.org/10.1145/3098243.3098264

Malina, L., Hajny, J., Fujdiak, R., & Hosek, J. (2016). On perspective of security and privacy-preserving solutions in the internet of things. *Computer Networks*, *102*, 83–95. https://doi.org/10.1016/j.comnet.2016.03.011

Maple, C. (2017). Security and privacy in the internet of things. *Journal of Cyber Policy*, *2*(2), 155–184. https://doi.org/10.1080/23738871.2017.1366536

Menard, P., Bott, G. J., & Crossler, R. E. (2017). User Motivations in Protecting Information Security: Protection Motivation Theory Versus Self-Determination Theory. *Journal of Management Information Systems*, *34*(4), 1203–1230. https://doi.org/10.1080/07421222.2017.1394083

Michie, S., van Stralen, M. M., & West, R. (2011). The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science*, *6*(1), 42. https://doi.org/10.1186/1748-5908-6-42

Mirkovic, J., & Reiher, P. (2004). A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review*, *34*(2), 39–53. https://doi.org/10.1145/997150.997156

Mohana Priya, P., Akilandeswari, V., Mercy Shalinie, S., Lavanya, V., & Shanmuga Priya, M. (2014). The Protocol Independent Detection and Classification (PIDC) system for DRDoS attack. *2014 International Conference on Recent Trends in Information Technology*, 1–7. https://doi.org/10.1109/ICRTIT.2014.6996154

Moore, T., Clayton, R., & Anderson, R. (2009). The Economics of Online Crime. *Journal of Economic Perspectives*, *23*(3), 3–20. https://doi.org/10.1257/jep.23.3.3

Nguyen, S. D., Mimura, M., & Tanaka, H. (2018). Abusing TCP Retransmission for DoS Attack Inside Virtual Network. In *International workshop on information security applications* (pp. 199–211). https://doi.org/10.1007/978-3-319-93563-8{_}17

Nosyk, Y., Korczyński, M., & Duda, A. (2022). Routing Loops as Mega Amplifiers for DNS-Based DDoS Attacks. In *International conference on passive and active network measurement* (pp. 629–644). https://doi.org/10.1007/978-3-030-98785-5{_}28

Nuiaa, R. R., Manickam, S., & Alsaeedi, A. H. (2021). Distributed reflection denial of service attack: A critical review. *International Journal of Electrical and Computer Engineering (IJECE)*, *11*(6), 5327. https://doi.org/10.11591/ijece.v11i6.pp5327-5341

Omolara, A. E., Alabdulatif, A., Abiodun, O. I., Alawida, M., Alabdulatif, A., Alshoura, W. H., & Arshad, H. (2022). The internet of things security: A survey encompassing unexplored areas and new insights. *Computers & Security*, *112*, 102494. https://doi.org/10.1016/j.cose.2021.102494

Ormond, D., & Barlow, J. (2022). Security Warning Messages Research: Past and Future. *MWAIS 2022 Proceedings*.

Pal, D. (2022). A new DDoS attack vector: TCP Middlebox Reflection. https://blog.apnic.net/2022/10/18/a-new-ddos-attack-vector-tcp-middlebox-reflection/

Palo Alto Networks. (2022). CVE-2022-0028 PAN-OS: Reflected Amplification Denial-of-Service (DoS) Vulnerability in URL Filtering. https://security.paloaltonetworks.com/CVE-2022-0028

Panahi Rizi, M. H., & Hosseini Seno, S. A. (2022). A systematic review of technologies and solutions to improve security and privacy protection of citizens in the smart city. *Internet of Things*, *20*, 100584. https://doi.org/10.1016/j.iot.2022.100584

Panchiwala, S., & Shah, M. (2020). A Comprehensive Study on Critical Security Issues and Challenges of the IoT World. *Journal of Data, Information and Management*, *2*(4), 257–278. https://doi.org/10.1007/s42488-020-00030-2

Prajapati, S., & Singh, A. (2022). Cyber-Attacks on Internet of Things (IoT) Devices, Attack Vectors, and Remedies: A Position Paper. https://doi.org/10.1007/978-3-030-73885-3{_}17

Rajivan, P., Aharonov-Majar, E., & Gonzalez, C. (2020). Update now or later? Effects of experience, cost, and risk preference on update decisions. *Journal of Cybersecurity*, *6*(1). https://doi.org/10.1093/cybsec/tyaa002

Rijksoverheid. (2021). Hou tv en deurbel slim: Doe je updates. https://www.rijksoverheid.nl/ministeries/ministerie-van-economische-zaken-en-klimaat/nieuws/2021/12/13/hou-tv-en-deurbel-slim-doe-je-updates

Rizvi, S., Orr, R., Cox, A., Ashokkumar, P., & Rizvi, M. R. (2020). Identifying the attack surface for IoT network. *Internet of Things*, *9*, 100162. https://doi.org/10.1016/j.iot.2020.100162

Rodriguez, E., Fukkink, M., Parkin, S., van Eeten, M., & Ganan, C. (2022). Difficult for Thee, But Not for Me: Measuring the Difficulty and User Experience of Remediating Persistent IoT Malware. *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*, 392–409. https://doi.org/10.1109/EuroSP53844.2022.00032

Rodríguez, E., Verstegen, S., Noroozian, A., Inoue, D., Kasama, T., van Eeten, M., & Gañán, C. H. (2021). User compliance and remediation success after IoT malware notifications. *Journal of Cybersecurity*, *7*(1). https://doi.org/10.1093/cybsec/tyab015

Rogers, R. W. (1975). A Protection Motivation Theory of Fear Appeals and Attitude Change1. *The Journal of Psychology*, *91*(1), 93–114. https://doi.org/10.1080/00223980.1975.9915803

Rogers, R. W. (1983). Cognitive and psychological processes in fear appeals and attitude change: A revised theory of protection motivation. In *Social psychophysiology: A sourcebook* (pp. 153–176).

Rossow, C. (2014). Amplification Hell: Revisiting Network Protocols for DDoS Abuse. *Proceedings 2014 Network and Distributed System Security Symposium*. https://doi.org/10.14722/ndss.2014.23233

Rowe, B., Wood, D., Reeves, D., & Braun, F. (2011). The Role of Internet Service Providers in Cyber Security. *Institute for Homeland Security Solutions*.

Sarabi, A., Zhu, Z., Xiao, C., Liu, M., & Dumitraş, T. (2017). Patch Me If You Can: A Study on the Effects of Individual User Behavior on the End-Host Vulnerability State. In *Pam 2017: Passive and active measurement* (pp. 113–125). https://doi.org/10.1007/978-3-319-54328-4{_}9

Sen, R., Choobineh, J., & Kumar, S. (2020). Determinants of Software Vulnerability Disclosure Timing. *Production and Operations Management*, *29*(11), 2532–2552. https://doi.org/10.1111/poms.13120

Shadowserver. (2022). Vulnerable DDoS Middlebox Report. https://www.shadowserver.org/what-we-do/network-reporting/vulnerable-ddos-middlebox-report/

Shadowserver. (2023). Who We Are. https://www.shadowserver.org/who-we-are/

Sicari, S., Rizzardi, A., Grieco, L., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, *76*, 146–164. https://doi.org/10.1016/j.comnet.2014.11.008

Simpson, A. K., Roesner, F., & Kohno, T. (2017). Securing vulnerable home IoT devices with an in-hub security manager. *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, 551–556. https://doi.org/10.1109/PERCOMW.2017.7917622

Sinanovic, H., & Mrdovic, S. (2017). Analysis of Mirai malicious software. *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, 1–5. https://doi.org/10.23919/SOFTCOM.2017.8115504

Stock, B., Pellegrino, G., Backes, M., & Rossow, C. (2018). Didn't You Hear Me? - Towards More Successful Web Vulnerability Notifications. *Proceedings 2018 Network and Distributed System Security Symposium*. https://doi.org/10.14722/ndss.2018.23171

Stock, B., Pellegrino, G., Rossow, C., Johns, M., & Backes, M. (2016). Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. *USENIX Security Symposium (USENIX Security 16)*, 1015–1032. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock

Tabassum, M., Kosinski, T., & Lipford, H. R. (2019). I don't own the data": End User Perceptions of Smart Home Device Data Practices and Risks. *SOUPS@ USENIX Security Symposium*.

Terrell, S. R. (2012). Mixed-Methods Research Methodologies. *The Qualitative Report*, *17*(1), 254–280.

Transforma Insights. (2022). Number of Internet of Things (IoT) connected devices worldwide from 2019 to 2021, with forecasts from 2022 to 2030. https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/

Tsai, H.-y. S., Jiang, M., Alhabash, S., LaRose, R., Rifon, N. J., & Cotten, S. R. (2016). Understanding online safety behaviors: A protection motivation theory perspective. *Computers & Security*, *59*, 138–150. https://doi.org/10.1016/j.cose.2016.02.009

UN. (2023). Cybersecurity. https://www.un.org/counterterrorism/cybersecurity

van der Kleij, R., van 't Hoff-De Goede, S., van de Weijer, S., & Leukfeldt, R. (2021). How Safely Do We Behave Online? An Explanatory Study into the Cybersecurity Behaviors of Dutch Citizens. https://doi.org/10.1007/978-3-030-79997-7{_}30

van Eeten, M. (2017). Patching security governance: an empirical view of emergent governance mechanisms for cybersecurity. *Digital Policy, Regulation and Governance*, *19*(6), 429–448. https://doi.org/10.1108/DPRG-05-2017-0029

Wang, Z., Zhang, Y., & Liu, Q. (2012). A research on vulnerability discovering for router protocols based on fuzzing. *7th International Conference on Communications and Networking in China*, 245–250. https://doi.org/10.1109/ChinaCom.2012.6417484

West, R., & Michie, S. (2020). A brief introduction to the COM-B Model of behaviour and the PRIME Theory of motivation. *Qeios*. https://doi.org/10.32388/WW04E6.2

Wogalter, M. S. (2006). Communication-human information processing (C-HIP) model. In *Handbook of warnings* (pp. 51–61). Lawrence Erlbaum Associates Mahwah, NJ.

Yang, Y., Wu, L., Yin, G., Li, L., & Zhao, H. (2017). A Survey on Security and Privacy Issues in Internet-of-Things. *IEEE Internet of Things Journal*, *4*(5), 1250–1258. https://doi.org/10.1109/JIOT.2017.2694844

Zeng, E., Li, F., Stark, E., Porter, A., & Tabriz, F. P. (2019). Fixing HTTPS Misconfigurations at Scale: An Experiment with Security Notifications. *The 2019 Workshop on the Economics of Information Security*.

Zhang, J., Duan, H., Liu, W., & Yao, X. (2017). How to notify a vulnerability to the right person? Case study: in an ISP scope. *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, 1–7.

Zhang, Z.-K., Cho, M. C. Y., Wang, C.-W., Hsu, C.-W., Chen, C.-K., & Shieh, S. (2014). IoT Security: Ongoing Challenges and Research Opportunities. *2014 IEEE 7th International Conference on Service-Oriented Computing and Applications*, 230–234. https://doi.org/10.1109/SOCA.2014.58

A

Literature study

To identify relevant academic literature, which is presented in chapter 2, a boolean search strategy is used. This appendix provides an overview of the used keywords and combinations per section of the theoretical background chapter. For finding literature, mainly the Scopus databases are used. Moreover, to complement the findings from Scopus, Google Scholar is used. However, it is tried to minimise the use of Google Scholar because it is unclear how articles are ranked in this database (Beel et al., 2009). Table 10 shows the different keyword combinations used for performing the literature study.

Table 10: Boolean keywords literature study

Topic	**Keywords**
Security issues of internet-connected devices	(IoT OR internet of things) AND security AND challenges
	(IoT OR internet of things) AND (security OR access control OR authentication challenges)
DDoS attacks	DDoS AND (attack OR attack types)
	(UDP OR TCP) AND reflection AND attacks
	TCP AND reflective AND amplification AND attack
	TCP AND amplification AND attack
Vulnerability notifications	Vulnerability AND notification
	Abuse AND notification AND vulnerable AND IoT
End-user Security behaviour	IoT AND patching
	Cybersecurity AND (warning OR notification) AND consumer behaviour
	com-b AND behaviour AND model AND (security OR cybersecurity)
	C-HIP AND model AND cybersecurity
	user AND security AND behaviour AND motivations
	vulnerability AND lifecycle AND patching
	(end-user OR user) AND security AND patching AND behaviour

B

Python script custom TCP packets

This appendix contains the python code used for sending the custom TCP packets. On line 4, the IP address of the destination is defined. As this code serves as an example, no IP address is defined and solely the format of an IP address is shown.

```
1  from scapy.all import *
2  import random
3
4  dst = 'xx.xx.xx.xx'
5  sport=random.randint(1024,65533)
6  dport = 80
7  seq=random.randint(1000,4294967290)
8  ipid=random.randint(1024,65533)
9
10 #SYN with GET
11 packet = IP(dst=dst, id=ipid) / TCP(sport=sport, dport=dport, flags='S', seq=seq)
       / "GET / HTTP/#1.1\r\nHost: www.kpn.com\r\n\r\n"
12 print("sending single SYN with srcport %s to destip %s" % (str(sport),str(dst)))
13 send(packet)
14 SR = sr1(packet)
15 SR.show()
16 print(SR[TCP].payload)
17
18 #SYN + PUSH | ACK with GET
19 sport=random.randint(1024,65533)
20 seq=random.randint(1000,4294967290)
21 ack=random.randint(1000,4294967290)
22 ipid=random.randint(1024,65533)
23 packet1 = IP(dst=dst, id=ipid) / TCP(sport=sport, dport=dport, flags='S', seq=seq)
24 packet2 = IP(dst=dst, id=ipid) / TCP(sport=sport, dport=dport, flags='PA', seq=seq
       +1, ack=ack) / "GET / HTTP/1.1\r\nHost: www.kpn.com\r\n\r\n"
25 print("sending SYN + PUSH & ACK with GET with srcport %s" % (str(sport)))
26 send(packet1)
27 send(packet2)
```

Listing B.1: Python code for custom TCP packets

C

Vulnerability notifications

This appendix contains the vulnerability notification sent to customers in TG1. All the notifications are shown in Dutch, as there was no English version created and send to customers. This mail can be seen in the figures 19 and 20. Furthermore, the email sent to customers in TG2 and TG3 is displayed in figures 21 and 22. Last, in figure 23 the web page on the KPN website is shown which provided the contacted customers additional information on the vulnerability and the research.

Oortwijn, Joost

From:	KPN Abuse Team <abuse@kpn.com>
Sent:	maandag 27 februari 2023 16:04
To:	Oortwijn, Joost
Subject:	Misbruik van uw Internetverbinding j.oortwijn02@kpn.com
Attachments:	KPN Hussel logo; uitroepteken_driehoek; Shadow 1; Logo footer; Privacy waarborg; icon-arrow.png

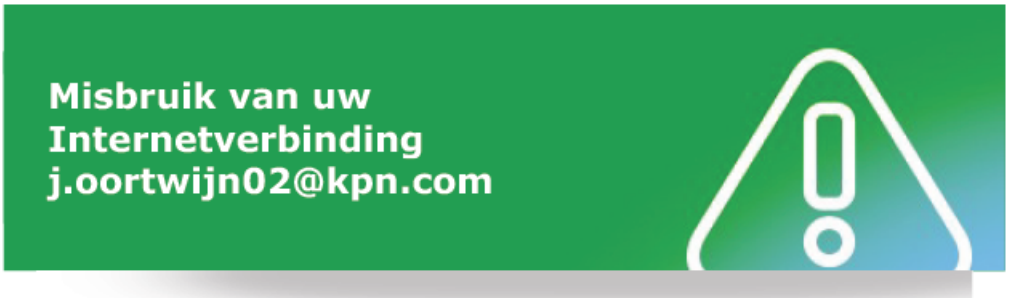

Beste meneer/mevrouw ,

Veilig internet is voor iedereen belangrijk. Daarom houden we samen onze internetverbindingen zo veilig mogelijk. We hebben hierbij uw hulp nodig, omdat we een veiligheidsprobleem hebben ontdekt op uw internetverbinding. Doorloop alstublieft de stappen hieronder vandaag nog.

Waarom moet ik iets doen?
We lossen de meeste veiligheidsproblemen op afstand op. Maar soms hebben we de hulp van onze klanten nodig. Het is voor u en andere klanten belangrijk dat u helpt, ook als u zelf niets merkt van het veiligheidsprobleem.

Wat is er aan de hand en hoe kan ik dit oplossen?
Een apparaat die gebruik maakt van uw internetverbinding kan door kwaadwillenden misbruikt worden voor het uitvoeren van grootschalige aanvallen op het internet. Op dit moment kunnen wij niet met zekerheid vaststellen om welk apparaat het gaat. Het meest aannemelijk is dat het een apparaat is dat onderdeel is van uw netwerk infrastructuur.

Voorbeelden van netwerk infrastructuur die kwetsbaar kunnen zijn voor misbruik:
- Firewalls
- Routers/modems

Daarnaast zou de kwetsbaarheid op uw internet verbinding veroorzaakt kunnen

3

Figure 19: Vulnerability notification TG1 (1)

worden door een van de volgende (slimme) apparaten.

Voorbeelden van slimme apparaten die kwetsbaar zijn voor misbruik:
• Energie monitoring systemen
• Alarm systemen
• (Beveiliging) camera's
• Woning beheer systemen (temperatuur, licht, geluid of gordijn bediening)

De kwetsbaarheid van uw apparaat wordt vaak veroorzaakt door een verkeerde configuratie door de fabrikant. Hier kunt u zelf vaak weinig aan doen. De fabrikant brengt updates uit om uw apparaat veilig te houden. Wij adviseren u daarom om uw (slimme) apparaten altijd up-to-date te houden met de meeste recente software die door de fabrikant wordt uitgebracht. In het bijzonder vragen wij u om voor de bovengenoemde apparaten de meest actuele software/update te installeren. Indien uw apparaat de mogelijkheid heeft om een wachtwoord in te stellen raden we aan om van deze mogelijkheid gebruik te maken en een sterk en uniek wachtwoord in te stellen.

Voor meer informatie zie https://kpn.com/onderzoek-tudelft

De afdeling Abuse
De afdeling Abuse handelt veiligheidsincidenten af voor KPN.

Meer informatie

Hebt u nog vragen?
U kunt uw vragen stellen via e-mail op abuse@kpn.com.

Met vriendelijke groet,

KPN Abuse Team

Wat vindt u van deze e-mail?

Heel goed Kan beter

kp KPN B.V. - Postbus 30000 - 2500 GA Den Haag - KvK nr. 27124701 PRIVACY

4

Figure 20: Vulnerability notification TG1 (2)

Oortwijn, Joost

From:	KPN Abuse Team <abuse@kpn.com>
Sent:	maandag 27 februari 2023 16:05
To:	Oortwijn, Joost
Subject:	Beveiligingsmelding j.oortwijn02@kpn.com
Attachments:	KPN Hussel logo; uitroepteken_driehoek; Shadow 1; Logo footer; Privacy waarborg; icon-arrow.png

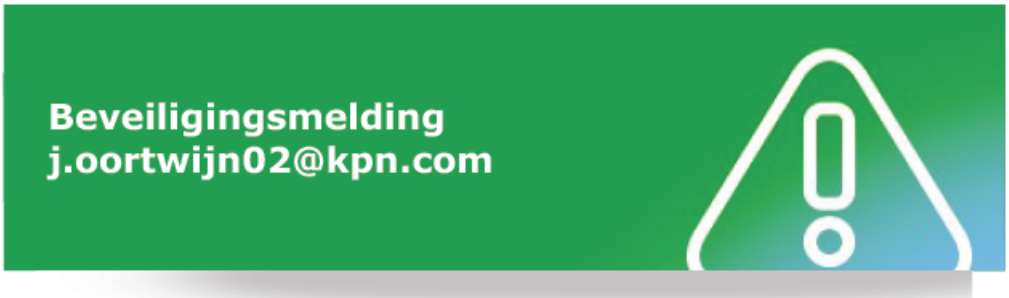

Beste meneer/mevrouw ,

Veilig internet is voor iedereen belangrijk. Daarom houden we samen onze internetverbindingen zo veilig mogelijk. We hebben hierbij uw hulp nodig, omdat we een veiligheidsprobleem hebben ontdekt op uw internetverbinding. Doorloop alstublieft de stappen hieronder vandaag nog.

Waarom moet ik iets doen?
We lossen de meeste veiligheidsproblemen op afstand op. Maar soms hebben we de hulp van onze klanten nodig. Het is voor u en andere klanten belangrijk dat u helpt, ook als u zelf niets merkt van het veiligheidsprobleem.

Wat is er aan de hand en hoe kan ik dit oplossen?
U ontvangt deze mail omdat wij recent, in de afgelopen zes maanden, een mogelijk beveiliging probleem hebben waargenomen op uw internetverbinding. Een apparaat die gebruik maakt van uw internetaansluiting kon door kwaadwillenden worden misbruikt voor het versturen van grootschalige aanvallen op het internet. Het ging hierbij om een (slim) apparaat dat zelfstandig gebruik maakte van uw internetverbinding.

In het kader van ons onderzoek naar dit beveiligingsprobleem zal onze collega Joost Oortwijn donderdag 2 februari telefonisch contact met u opnemen om u hierover enkele vragen te stellen, dit duurt ongeveer 5 minuten.

Onze gegevens laten zien dat deze mogelijke kwetsbaarheid op uw internet aansluiting

Figure 21: Email TG2 & TG3 (1)

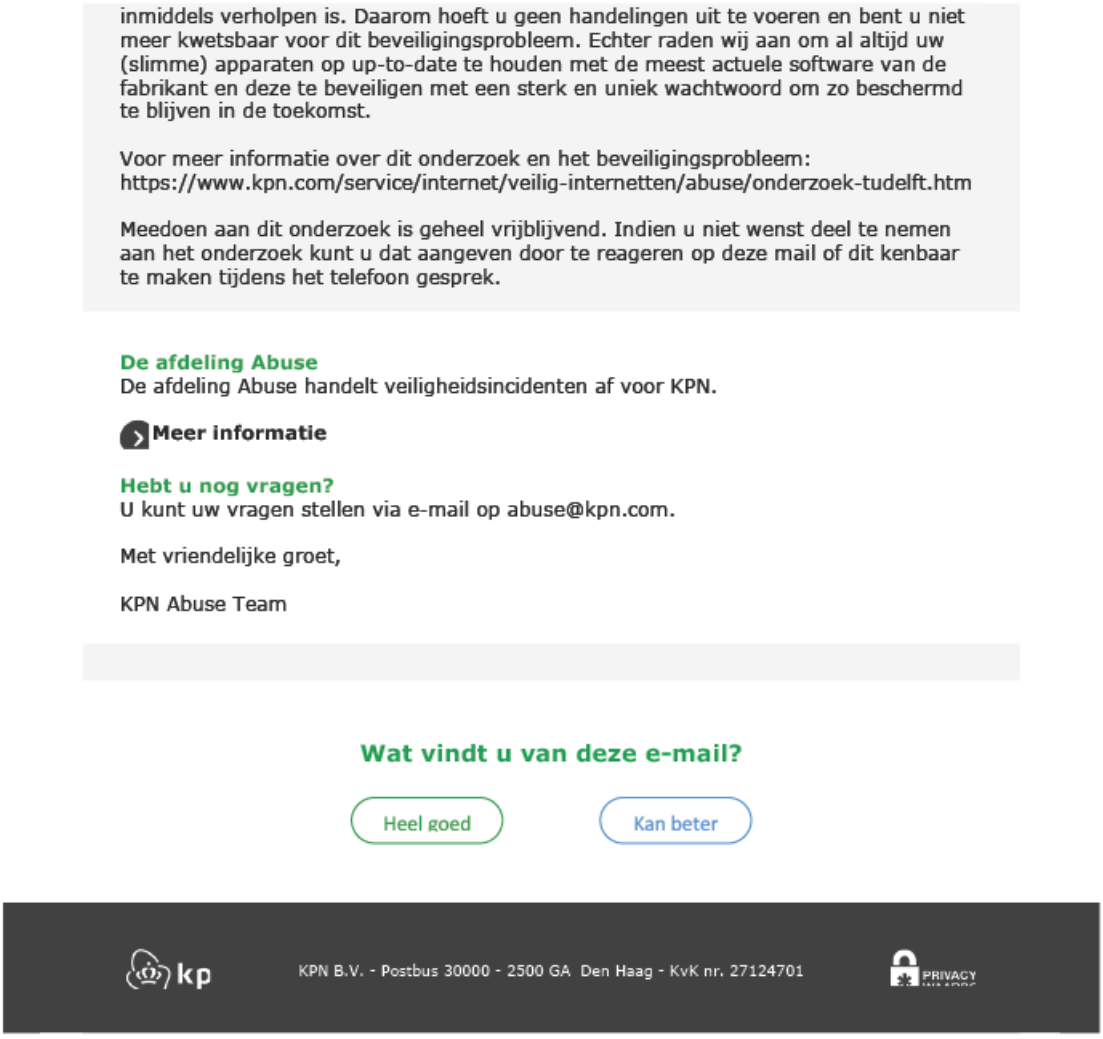

inmiddels verholpen is. Daarom hoeft u geen handelingen uit te voeren en bent u niet meer kwetsbaar voor dit beveiligingsprobleem. Echter raden wij aan om al altijd uw (slimme) apparaten op up-to-date te houden met de meest actuele software van de fabrikant en deze te beveiligen met een sterk en uniek wachtwoord om zo beschermd te blijven in de toekomst.

Voor meer informatie over dit onderzoek en het beveiligingsprobleem:
https://www.kpn.com/service/internet/veilig-internetten/abuse/onderzoek-tudelft.htm

Meedoen aan dit onderzoek is geheel vrijblijvend. Indien u niet wenst deel te nemen aan het onderzoek kunt u dat aangeven door te reageren op deze mail of dit kenbaar te maken tijdens het telefoon gesprek.

De afdeling Abuse
De afdeling Abuse handelt veiligheidsincidenten af voor KPN.

Meer informatie

Hebt u nog vragen?
U kunt uw vragen stellen via e-mail op abuse@kpn.com.

Met vriendelijke groet,

KPN Abuse Team

Wat vindt u van deze e-mail?

Heel goed Kan beter

kp KPN B.V. - Postbus 30000 - 2500 GA Den Haag - KvK nr. 27124701 PRIVACY

2

Figure 22: Email TG2 & TG3 (2)

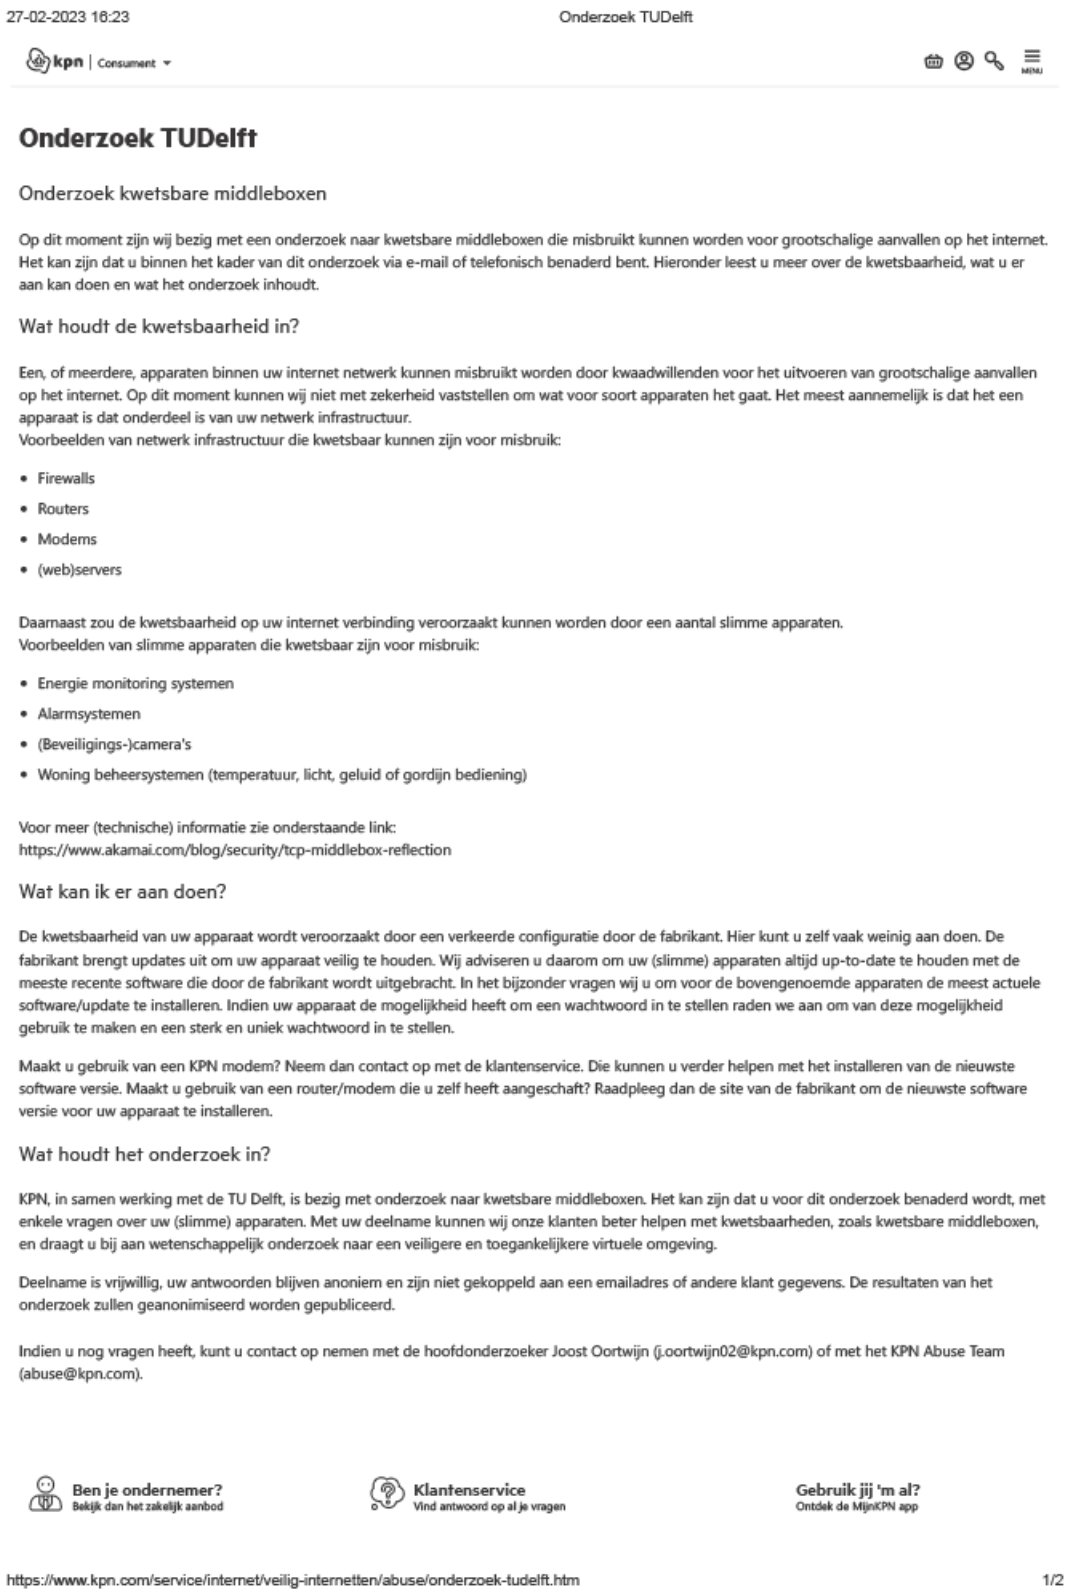

27-02-2023 16:23 Onderzoek TUDelft

kpn | Consument MENU

Onderzoek TUDelft

Onderzoek kwetsbare middleboxen

Op dit moment zijn wij bezig met een onderzoek naar kwetsbare middleboxen die misbruikt kunnen worden voor grootschalige aanvallen op het internet. Het kan zijn dat u binnen het kader van dit onderzoek via e-mail of telefonisch benaderd bent. Hieronder leest u meer over de kwetsbaarheid, wat u er aan kan doen en wat het onderzoek inhoudt.

Wat houdt de kwetsbaarheid in?

Een, of meerdere, apparaten binnen uw internet netwerk kunnen misbruikt worden door kwaadwillenden voor het uitvoeren van grootschalige aanvallen op het internet. Op dit moment kunnen wij niet met zekerheid vaststellen om wat voor soort apparaten het gaat. Het meest aannemelijk is dat het een apparaat is dat onderdeel is van uw netwerk infrastructuur.
Voorbeelden van netwerk infrastructuur die kwetsbaar kunnen zijn voor misbruik:

- Firewalls
- Routers
- Modems
- (web)servers

Daarnaast zou de kwetsbaarheid op uw internet verbinding veroorzaakt kunnen worden door een aantal slimme apparaten.
Voorbeelden van slimme apparaten die kwetsbaar zijn voor misbruik:

- Energie monitoring systemen
- Alarmsystemen
- (Beveiligings-)camera's
- Woning beheersystemen (temperatuur, licht, geluid of gordijn bediening)

Voor meer (technische) informatie zie onderstaande link:
https://www.akamai.com/blog/security/tcp-middlebox-reflection

Wat kan ik er aan doen?

De kwetsbaarheid van uw apparaat wordt veroorzaakt door een verkeerde configuratie door de fabrikant. Hier kunt u zelf vaak weinig aan doen. De fabrikant brengt updates uit om uw apparaat veilig te houden. Wij adviseren u daarom om uw (slimme) apparaten altijd up-to-date te houden met de meeste recente software die door de fabrikant wordt uitgebracht. In het bijzonder vragen wij u om voor de bovengenoemde apparaten de meest actuele software/update te installeren. Indien uw apparaat de mogelijkheid heeft om een wachtwoord in te stellen raden we aan om van deze mogelijkheid gebruik te maken en een sterk en uniek wachtwoord in te stellen.

Maakt u gebruik van een KPN modem? Neem dan contact op met de klantenservice. Die kunnen u verder helpen met het installeren van de nieuwste software versie. Maakt u gebruik van een router/modem die u zelf heeft aangeschaft? Raadpleeg dan de site van de fabrikant om de nieuwste software versie voor uw apparaat te installeren.

Wat houdt het onderzoek in?

KPN, in samen werking met de TU Delft, is bezig met onderzoek naar kwetsbare middleboxen. Het kan zijn dat u voor dit onderzoek benaderd wordt, met enkele vragen over uw (slimme) apparaten. Met uw deelname kunnen wij onze klanten beter helpen met kwetsbaarheden, zoals kwetsbare middleboxen, en draagt u bij aan wetenschappelijk onderzoek naar een veiligere en toegankelijkere virtuele omgeving.

Deelname is vrijwillig, uw antwoorden blijven anoniem en zijn niet gekoppeld aan een emailadres of andere klant gegevens. De resultaten van het onderzoek zullen geanonimiseerd worden gepubliceerd.

Indien u nog vragen heeft, kunt u contact op nemen met de hoofdonderzoeker Joost Oortwijn (j.oortwijn02@kpn.com) of met het KPN Abuse Team (abuse@kpn.com).

Ben je ondernemer? Bekijk dan het zakelijk aanbod

Klantenservice Vind antwoord op al je vragen

Gebruik jij 'm al? Ontdek de MijnKPN app

https://www.kpn.com/service/internet/veilig-internetten/abuse/onderzoek-tudelft.htm 1/2

Figure 23: Information on the vulnerability and research on KPN website

D

End-user interview protocol

This appendix shows the different interview protocols used for the end-user interviews. In figure 24 and 25 part 1 and part 2 of the interview protocol for end-users in TG1 is shown. In figure 26 and 27 parts 1 and part 2 of the interview protocol for end-users in TG2 and TG3 is presented.

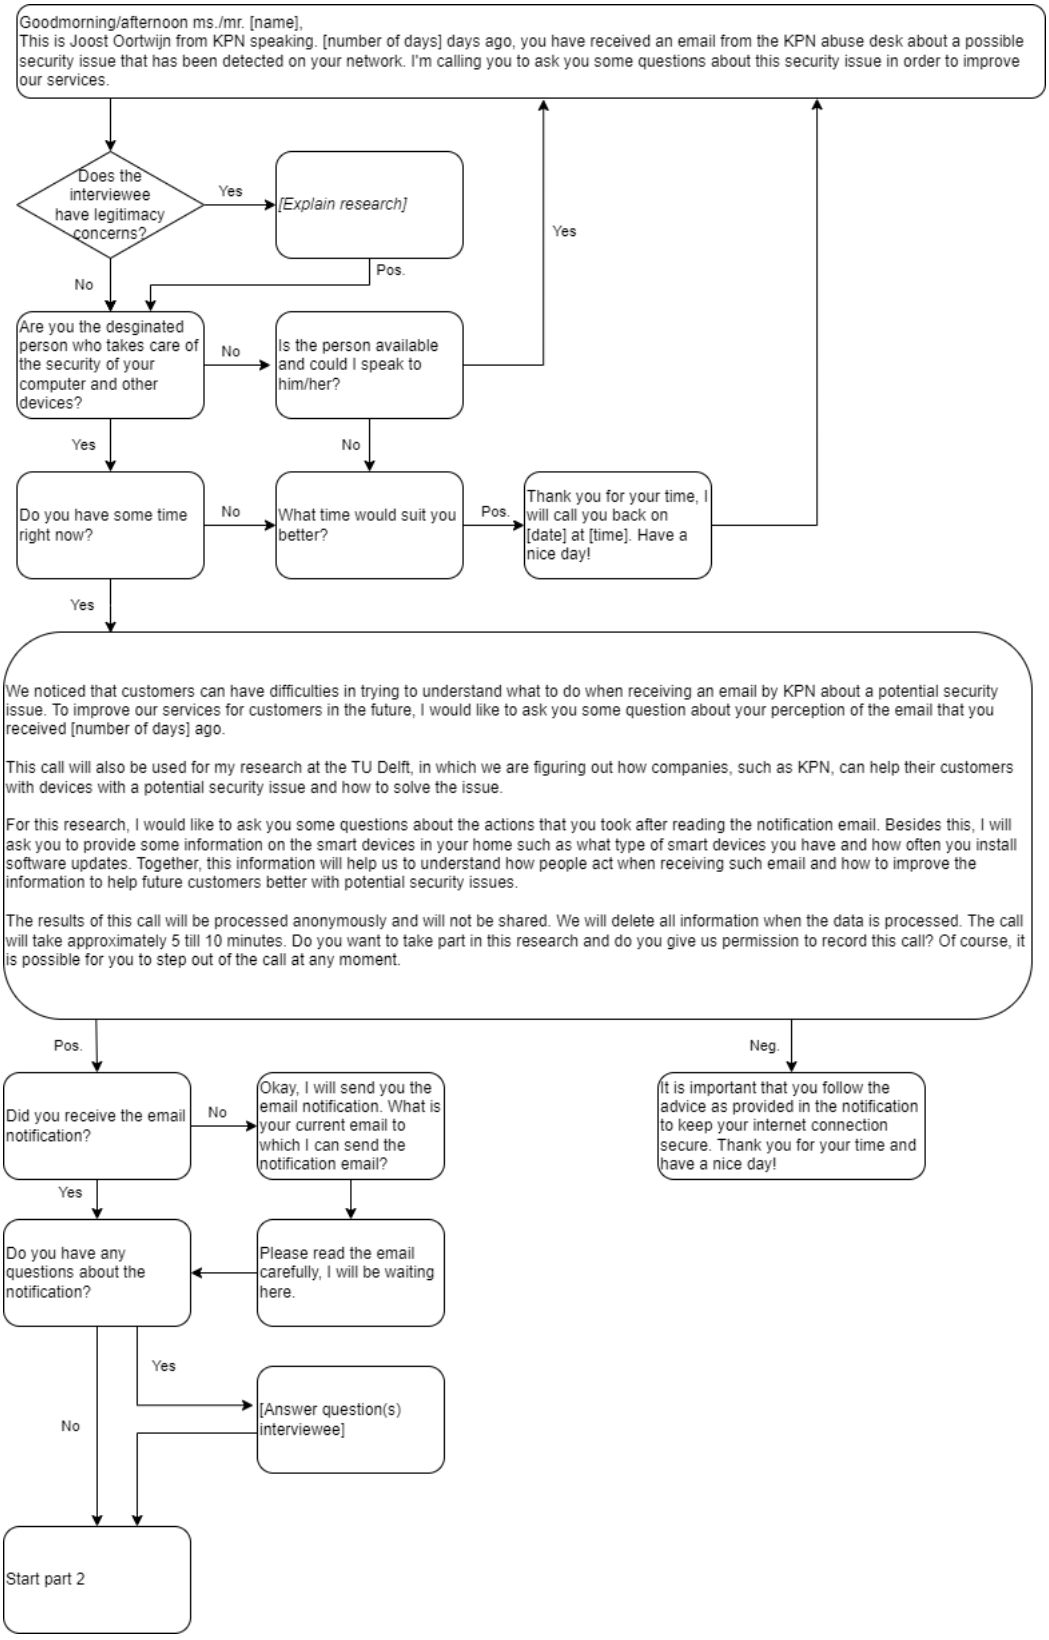

Figure 24: Interview protocol TG1 part 1

Figure 25: Interview protocol TG1 part 2

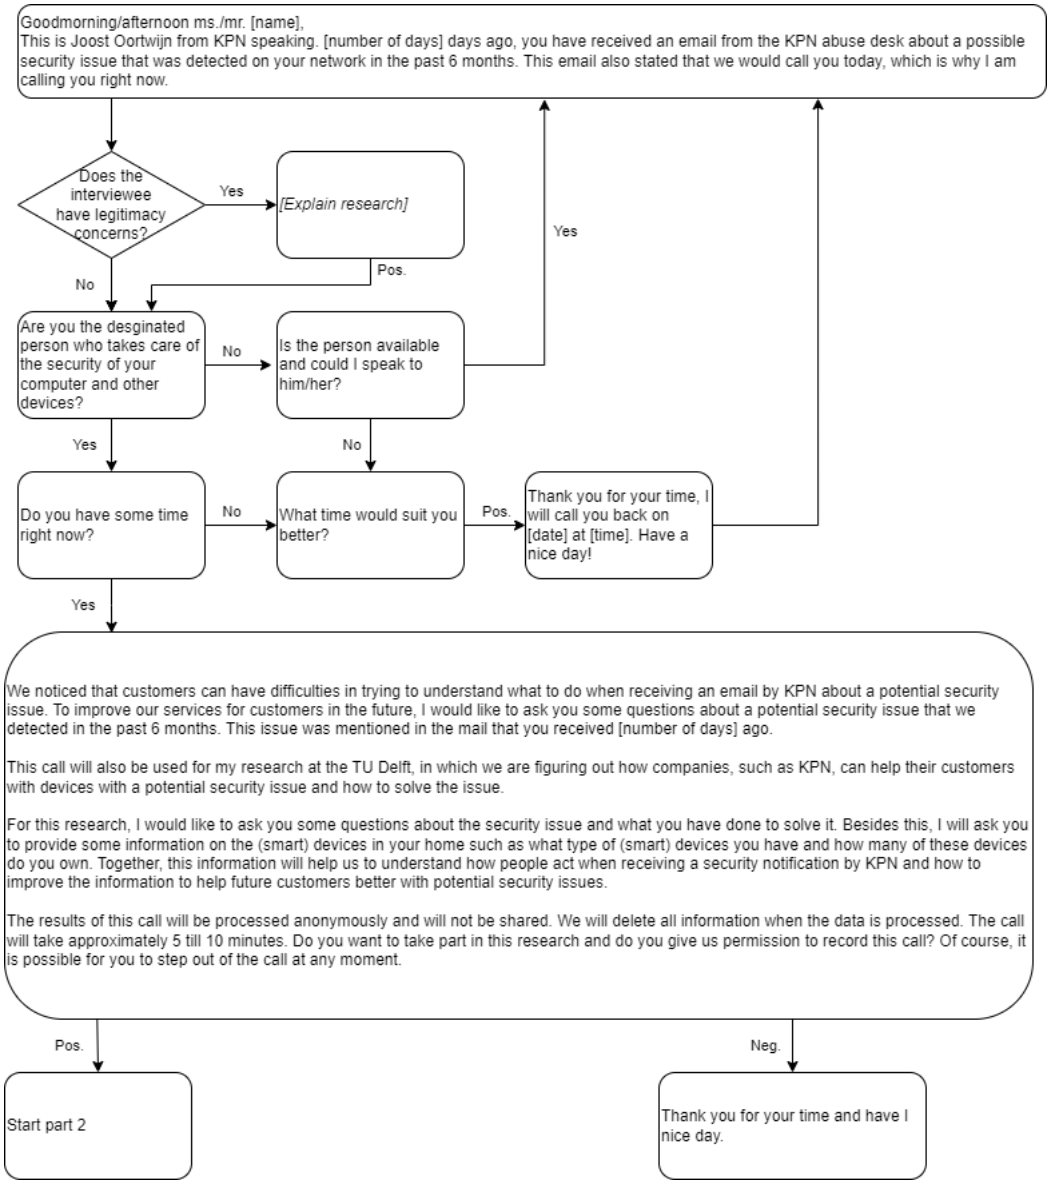

Figure 26: Interview protocol TG2 & TG3 part 1

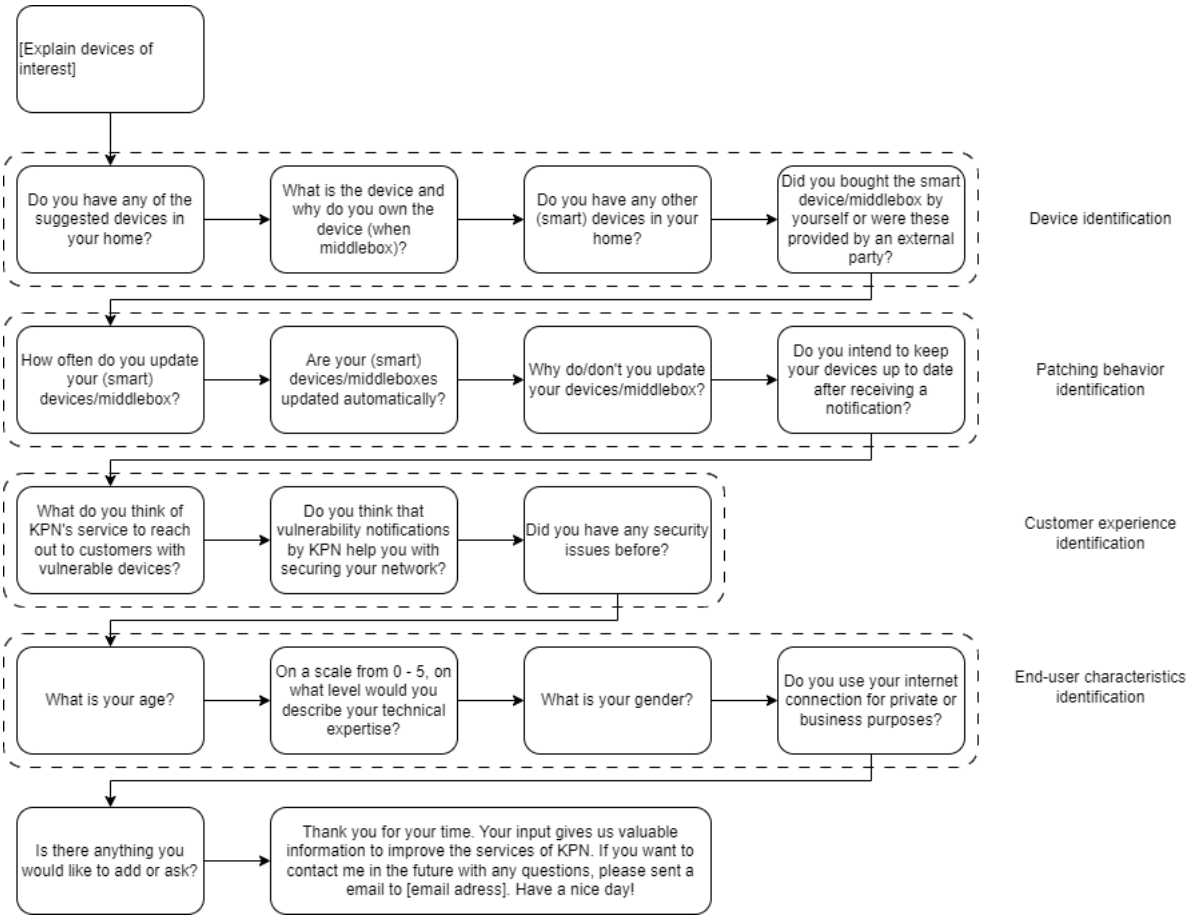

Figure 27: Interview protocol TG2 & TG3 part 2

E

Manufacturer interview protocol

This appendix shows the interview protocol which is used for interviewing the manufacturer in figure 28. The name of the interviewee, brand name and device name are anonymized for privacy purposes.

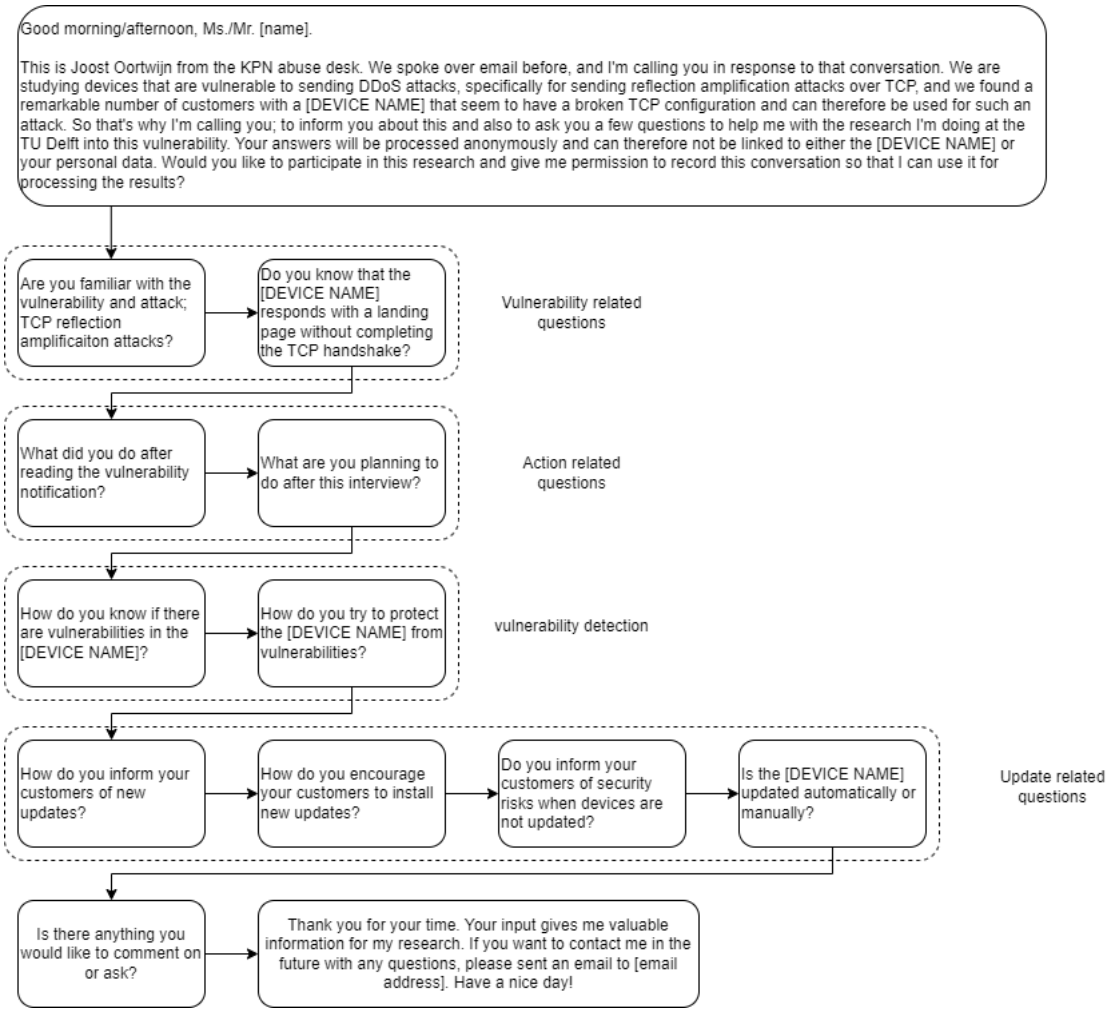

Figure 28: Interview protocol manufacturer

F

Descriptive Statistics

This appendix provides an overview of the descriptive statistics. This overview is presented below in table 11. The initial column indicates whether the corresponding customer was sent a vulnerability or email notification and contacted for the interview. However, this does not necessarily imply that the customer participated in the research, as there may be several reasons why a contacted individual may choose not to participate.

This table provides the output of the processed Shadowserver data, as discussed in chapter 4.2. However, the columns *"Mail"*, *"Target group"* and *"Device type"*, are manually added. The information in the column device type is identified by the various network scans as discussed in chapter 4.1.

Table 11: Overview descriptive statistics CM

	Mail	Target group	First seen	Last seen	Occurrences	Highest amp.	Ratio online	Device type
1	NO	TG3	2022-04-25	2022-08-10	68	159	63,6	NA
2	NO	TG3	2022-04-25	2022-08-11	71	78	65,7	NA
3	NO	TG3	2022-04-25	2022-08-10	23	8	21,5	NA
4	NO	TG3	2022-04-25	2022-08-11	64	78	59,3	NA
5	YES	TG1	2022-04-25	2023-02-19	192	1837	64	Energy monitor
6	NO	TG3	2022-04-25	2022-06-30	55	2	83,3	Firewall
7	NO	TG3	2022-04-25	2022-08-09	77	159	72,6	NA
8	NO	TG3	2022-04-25	2022-08-11	65	78	60,2	Router
9	NO	TG3	2022-04-25	2022-08-11	76	159	70,4	NA
10	NO	TG3	2022-04-25	2022-05-23	19	1	67,9	Firewall
11	NO	TG3	2022-04-25	2022-08-09	64	78	60,4	NA
12	YES	TG1	2022-04-25	2023-02-19	183	1756	61	Energy monitor
13	NO	TG3	2022-04-25	2022-04-25	1	1	100	Firewall
14	NO	TG3	2022-04-25	2022-08-11	56	159	51,9	Router
15	NO	TG3	2022-04-25	2022-08-10	66	8	61,7	(Web)-server
16	NO	TG3	2022-04-25	2022-06-28	44	8	68,8	NA
17	YES	TG1	2022-04-25	2023-02-18	103	251	34,4	BMS
18	NO	TG3	2022-04-25	2022-08-10	72	159	67,3	NA
19	NO	TG3	2022-04-25	2022-08-11	73	159	67,6	NA
20	YES	TG3	2022-04-27	2022-08-11	67	78	63,2	Firewall
21	NO	TG3	2022-04-27	2022-08-10	71	159	67,6	NA
22	NO	TG3	2022-04-27	2022-08-11	69	78	65,1	NA
23	YES	TG1	2022-04-27	2023-01-17	84	251	31,7	BMS
24	NO	TG3	2022-04-27	2022-05-01	4	24	100	NA
25	YES	TG3	2022-04-27	2022-08-11	63	78	59,4	Firewall
26	NO	TG3	2022-04-27	2022-04-27	1	1	100	BMS

Table 11 Continued from previous page

	Mail	Target group	First seen	Last seen	Occurrences	Highest amp.	Ratio online	Device type
27	NO	TG3	2022-04-27	2022-08-11	69	159	65,1	NA
28	NO	TG3	2022-04-27	2022-08-08	76	78	73,8	NA
29	NO	TG3	2022-04-27	2022-08-10	75	78	71,4	NA
30	NO	TG3	2022-04-27	2022-08-11	68	159	64,2	NA
31	NO	TG3	2022-04-28	2022-08-08	62	159	60,8	NA
32	YES	TG1	2022-04-28	2023-01-17	80	222	30,3	DVR
33	NO	TG3	2022-04-28	2022-08-09	50	159	48,5	DVR
34	NO	TG3	2022-04-28	2022-08-10	72	78	69,2	(Web)-server
35	YES	TG1	2022-04-28	2023-02-19	194	1307	65,3	Alarmsystem
36	NO	TG3	2022-04-28	2022-08-09	67	159	65	Firewall
37	YES	TG3	2022-04-28	2022-08-11	59	78	56,2	Firewall
38	NO	TG3	2022-04-28	2022-08-10	76	159	73,1	NA
39	NO	TG3	2022-04-28	2022-08-11	66	159	62,9	NA
40	NO	TG3	2022-04-29	2022-08-10	63	159	61,2	NA
41	YES	TG1	2022-04-29	2023-02-18	195	1837	66,1	Energy monitor
42	NO	TG3	2022-04-29	2022-08-11	62	159	59,6	NA
43	YES	TG1	2022-04-29	2023-02-19	195	71	65,9	Energy monitor
44	YES	TG3	2022-04-30	2022-08-11	75	159	72,8	Firewall
45	YES	TG3	2022-05-01	2022-08-10	75	159	74,3	Firewall
46	NO	TG3	2022-05-01	2022-08-10	69	159	68,3	Firewall
47	NO	TG3	2022-05-02	2022-06-05	9	1	26,5	Firewall
48	NO	TG3	2022-05-02	2022-08-08	6	8	6,1	Firewall
49	NO	TG3	2022-05-03	2022-08-11	52	15802	52	NA
50	NO	TG2	2022-05-03	2022-09-07	34	251	26,8	NA
51	NO	TG3	2022-05-06	2022-08-07	17	8	18,3	Firewall
52	NO	TG3	2022-05-06	2022-08-05	38	15798	41,8	NA
53	NO	TG3	2022-05-08	2022-08-30	47	1983	41,2	DVR
54	NO	TG3	2022-05-13	2022-08-11	56	8	62,2	Firewall
55	YES	TG1	2022-05-18	2023-02-16	181	71	66,1	Energy monitor
56	YES	TG1	2022-05-19	2023-01-17	160	1615	65,8	Energy monitor
57	NO	TG3	2022-05-23	2022-08-11	57	159	71,2	NA
58	NO	TG3	2022-06-05	2022-08-04	6	8	10	Firewall
59	YES	TG1	2022-07-01	2023-01-23	68	3539	33	BMS
60	YES	TG2	2022-07-09	2022-12-06	2	18,37	1,3	Energy monitor
61	NO	TG3	2022-07-10	2022-07-10	1	8	100	NA
62	NO	TG2	2022-07-15	2022-09-02	2	1956	4,1	NA
63	YES	TG3	2022-07-30	2022-08-14	15	150468	100	Firewall
64	YES	TG1	2022-08-05	2023-01-25	2	768	1,2	Energy monitor
65	YES	TG2	2022-08-11	2022-12-31	3	234	2,1	Router
66	YES	TG1	2022-08-19	2023-02-13	18	312	10,1	Washing machine
67	NO	TG3	2022-08-29	2022-08-29	1	1837	100	Energy monitor
68	NO	TG3	2022-08-30	2022-08-30	1	71	100	NA
69	NO	TG2	2022-10-13	2022-10-13	1	1412	100	NA
70	YES	TG2	2022-12-06	2022-12-06	1	1959	100	BMS
71	YES	TG2	2022-12-11	2022-12-11	1	1615	100	Energy monitor
72	YES	TG1	2022-12-13	2023-02-16	51	1256	78,5	Wifi extender
73	YES	TG2	2022-12-15	2022-12-15	1	1563	100	Energy monitor
74	YES	TG2	2022-12-30	2022-12-30	1	1837	100	Energy monitor
75	NO	TG2	2023-01-03	2023-01-03	1	366	100	NA
76	YES	TG1	2023-01-09	2023-01-23	8	1615	57,1	Energy monitor
77	YES	TG1	2023-01-30	2023-01-30	1	1615	100	Energy monitor
78	YES	TG1	2023-02-09	2023-02-19	10	1605	100	Energy monitor
79	YES	TG1	2023-02-10	2023-02-10	1	71	100	Energy monitor