DDOS MITIGATION

A MEASUREMENT-BASED APPROACH

Mattijs Jonker

DDoS Mitigation
A Measurement-Based Approach
Mattijs Jonker

Graduation Committee

Chairman: Prof. dr. J.N. Kok

Promotor: Prof. dr. ir. A. Pras

Co-promotor: Dr. A. Sperotto

Members:

Prof. dr.	K.C. Claffy	CAIDA, University of California, San Diego, USA
Prof. dr.	C. Rossow	CISPA, Saarland University, Germany
Prof. dr.	G. Smaragdakis	Technical University Berlin, Germany
Prof. dr.	J.L. van den Berg	University of Twente, The Netherlands
Prof. dr. ir.	L.J.M. Nieuwenhuis	University of Twente, The Netherlands

Funding sources

D3 - Distributed Denial-of-Service Defense – NWO project № 628.001.018

CONCORDIA – EU Horizon 2020 № 830927

Department of Homeland Security S&T/CSD № HHSP233201600012C

Air Force Research Laboratory № FA8750-18-2-0049

National Science Foundation № CNS-1730661 & ACI-1053575

DIGITAL SOCIETY INSTITUTE

DSI Ph.D. Thesis Series № 19-018

Digital Society Institute

P.O. Box 217

7500 AE Enschede, The Netherlands

ISBN: 978-90-365-4868-7

ISSN: 2589-7721 (DSI Ph.D. Thesis Series № 19-018)

DOI: 10.3990/1.9789036548687

<https://doi.org/10.3990/1.9789036548687>

Thesis typeset with L^AT_EX. Printed by Gildeprint, Enschede, The Netherlands.

Copyright © 2019 Mattijs Jonker – all rights reserved

No part of this publication may be reproduced without the prior written permission of the author.

DDOS MITIGATION: A MEASUREMENT-BASED APPROACH

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. T.T.M. Palstra,
volgens besluit van het College voor Promoties,
in het openbaar te verdedigen
op donderdag 10 oktober 2019 om 16:45

door

Mattijs Jonker

geboren op 30 januari 1983
te Alkmaar

This thesis has been approved by:

Prof. dr. ir. A. Pras (promotor)

Dr. A. Sperotto (co-promotor)

Acknowledgments

Dit proefschrift was niet tot stand gekomen zonder de kennis, hulp, vertrouwen en steun van anderen. Het is lastig om iedereen op een volstrekte wijze te bedanken. Toch richt ik me in dit dankwoord graag op een aantal mensen in het bijzonder.

Roland, je hebt mij in 2014, toen je een DNS meetvisie voor de toekomst had en zelf nog met je Ph.D. bezig was, de mogelijkheid geboden om aan *dnsjedi* te gaan samenwerken. Zonder dit project, wat beter als *OpenINTEL* bekend staat, had mijn proefschrift ongetwijfeld een andere invulling gekregen. Ik beschouw je als een briljante onderzoeker, voel me vereerd dat ik met je kan samenwerken, en bewonder de positieve invloed die je op anderen weet uit te oefenen. Ik ben ook blij dat je tijdens mijn verdediging als paranimf gereedstaat.

Aiko, bedankt voor je steun tijdens mijn promotietraject. Door noodlottige omstandigheden ben je niet altijd even nauw bij de technische aspecten van mijn werk betrokken geweest. Dat laat niet weg dat ik veel van je heb geleerd. Ik ben als persoon gegroeid onder jouw hoede en je daar dankbaar voor. Ik vind het bewonderingswaardig dat je zoveel wijsheid op mensen weet over te brengen en dat het menselijk vlak doorgaans prioriteit bij je krijgt.

Anna, bedankt dat je mij de kans hebt geboden om onder jouw supervisie aan mijn Ph.D. te gaan werken. Op academisch vlak heb je me veel bijgedragen. Onze ontelbare gesprekken waren essentieel voor de vorming van mijn proefschrift. En ik ben dankbaar voor het feit dat je, zelfs met twee jonge kinderen thuis, op sommige momenten 's avonds en zelfs tot diep in de nacht beschikbaar was om aan publicatie deadlines te werken.

Jeanette, bedankt dat je er zowel op persoonlijk als professioneel vlak voor me was over de afgelopen jaren. Jij bent een spilfiguur binnen onze vakgroep en de hoeveelheid werk die jij doet is werkelijk onbetaalbaar. Al mijn overige collega's binnen DACS wil ik bedanken voor de positieve werksfeer waaraan iedereen bijdraagt, alsmede de leuke momenten tijdens reizen naar conferenties, uitjes, enzovoort. Ik geloof dat een persoon kan bloeien mits omgeven door de juiste personen. Wat dat betreft voelt het alsof ik de loterij heb gewonnen.

Alberto Dainotti, I am deeply indebted to you for offering me the opportunity to come to CAIDA as visiting researcher. I believe my dissertation would not

have come to fruition without your bright insights and our collaboration. Thank you!

Alistair King, your expertise in systems engineering is something I admire. The amount of work you are able to do in relatively short time puts everybody else to shame. It was a pleasure to work with you. I also enjoyed our carpool rides at dawn and hope the electronic dance music I subjected you to in the car has not left your ears traumatized. *Josh Polterock*, you and I share an equal liking for coffee. And it was with one of many cups of the day that you always managed to spark the next conversation. Thank you for enriching my time abroad. To everybody else at *CAIDA*, thank you for being so welcoming of me during my stay abroad. It was a privilege to work in your midst and to be able to learn from a team that operates at the top.

To my graduation committee I wish to extend my sincere appreciation for taking the time to study my dissertation. I am honored to be considered worthy of your time. In particular, I would like to thank *Christian* and *Georgios* for traveling abroad to make it to my Ph.D. defence. And *kc* for attending remotely while recovering from an unfortunate bicycle accident.

Miriam, over de afgelopen jaren heb jij op belangrijke momenten voor me klaargestaan. Ik waardeer je vriendschap en ben blij dat je ook als paranimf voor me klaarstaat.

Thijs, je hebt niets meer mogen vernemen van mijn academische traject. Het zou je ongetwijfeld vreugde brengen en ik vind het erg jammer dat ik dit niet met je kan delen.

Finally, *Sabrina*, thank you for being there for me during the past years. Your close support has made a huge difference.

Mattijs

Abstract

Society heavily relies upon the Internet for global communications in this day and age. Although core Internet components were designed with resilience in mind, Internet stability and reliability are nowadays continuously subject to deliberate threats. These threats include Denial-of-Service (DoS) attacks, which can potentially be devastating.

About a decade ago, in 2010, the general public better started understanding the potential impact of DoS attacks. This was after a series of attacks by WikiLeaks supporters had caused pronounced disruption on the Internet. Various financial institutions were among the attacked targets, many of which saw their Web sites knocked offline or noticeably disrupted. A few years later, in 2013, the attack on *Spamhaus* shocked many with its record-breaking attack traffic volume. A few years after that, in 2016, a sequence of attacks on the DNS provider *Dyn* caused significant service outages, reverberating among large user bases in Europe and North America. The aforementioned examples may have a familiar ring as they are notorious cases. Yet they are ‘merely’ a selection of publicized incidents that underpin the gravity of the DoS threat. And while the DoS problem is by no means new, the number and intensity of attacks have especially over the past years reached unexpected proportions. In terms of sheer attack traffic volume, the bar is continually being raised. Think about, for example, recent reports of a 1.7 Tbps attack, which makes the once-shocking *Spamhaus* attack (300 Gbps) seem dinky in comparison. Experts argue that the full potential of attacks has not been seen yet, which prompts the question how many record-breaking attacks have yet to reach notoriety in the years to come.

As a result of attacks, not only businesses lose hundreds of millions of dollars annually. When it comes to vital infrastructure, national safety and even lives could be at stake. In the face of the evolving DoS threat, effective defenses are an absolute necessity. The upsurge of the DoS problem has prompted not only the development of diverse mitigation solutions, but has also given rise to a booming market for commercial products. Businesses and other prospective users of mitigation solutions find themselves having many shapes and sizes to choose from. The right fit may, however, not always be apparent. In addition, even though diverse solutions are readily available, their deployment and operation may come with hidden hazards that need to be better understood.

Policy makers and governments also find themselves facing questions concerning what needs to be done to promote cybersafety on a national level. Should we stimulate the market for mitigation solutions? Are there drawbacks to centralization of that market? And can we become too digitally dependent on other countries, especially when it comes to the safety and security of vital infrastructure? Given such questions, developing an optimal course of action to deal with the DoS problem brings about societal challenges that stack further upon technical ones.

Even though the DoS problem is not new, the scale of the problem is still unclear. We do not know exactly what it is we are defending against and getting a better understanding of attacks is essential to addressing the problem head-on. To advance situational awareness, many technical and societal challenges are yet to be tackled. Given the central importance of better understanding the DoS problem to improve overall Internet security, this thesis has three main contributions. First, this thesis rigorously characterizes DoS attacks and attacked targets at scale. Second, this thesis advances knowledge about the Internet-wide adoption, deployment and operational use of various mitigation solutions. Thirdly, this thesis investigates hidden hazards with mitigation solutions that have the potential to hamstring defenses or render mitigation solutions altogether ineffective.

In terms of the first contribution, this thesis reveals the massive scale of the DoS problem. To macroscopically characterize attacks and attack targets, we identify and systematically fuse diverse data from independent, global Internet measurement infrastructures. Our analysis of attacks reveals nearly 21 million attacks over a two-year period. We also show that, during the same period, about one-third of all /24 network address blocks estimated to be active on the Internet have been on the receiving end of at least one attack. This thesis also investigates the potential impact of attacks. We will show that Web hosting infrastructure is a prominent target and – using Web sites as a measure – we reveal that targeted infrastructure can be associated with well over 130 million Web sites (during a two-year period).

When it comes to the second contribution, this thesis investigates two solutions to mitigate attacks: cloud-based protection services and BGP blackholing. We quantify the uptake of protection services and reveal a prominent global trend in adoption. Our results highlight a relative growth in protection services use of $1.24\times$ (over a 1.5-year period) under the three top-level domains `com`, `net` and `org`, which combinedly account for about half of the global namespace. We also investigate the extent to which targets adopt (i.e., migrate to) protection services after having been targeted by a DoS attack. Our results highlight that attack intensity is an important factor for migration, whereas repeated attacks and attack duration are not. As for BGP blackholing, this thesis investigates

various operational aspects in the wild. Our results reveal how blackholing is applied in practice by operators. We show that for nearly 4% of attacks that are mitigated using blackholing, it takes more than 24 hours following the end of the attack for operators to retract the countermeasure. During this time, blackholed hosts are cutoff from the Internet (at least partially). The apparent lack of automation in recovery raises concern that hosts as well as services running on them may be cutoff from users unnecessarily. In addition, we unveil that less intense attacks are also blackholed: in 13% of cases the inferred attack traffic volume is at most 3 *Mbps*. As blackholing effectively brings about a ‘self-inflicted’ DoS, these findings raise the question of how much (or little) effort is required for attackers to get operators to trigger such an extreme countermeasure.

Focusing on the third contribution, this thesis investigates, for both mitigation solutions under consideration, hazards that can hamstring DoS defenses. Cloud-based protection services may be bypassed by sophisticated attackers as a result of mistakes in deployment. Mistakes may not be clearly understood by all users, which can lead to a false sense of security. We quantify this drawback on the Internet, focusing on the world’s most popular Web sites, and on leading commercial protection services. Our results underpin the extent of the problem: the protection of 41% of Web sites under consideration can be bypassed. As for blackholing, this thesis takes preliminary steps towards investigating the extent to which hosts are cutoff unnecessarily. We quantify this in terms of common Internet services that run on blackholed hosts.

This thesis will show from its outset that a basic challenge that we are faced with concerns data. Acquiring and developing diverse (raw) data sources to methodologically study the DoS problem constitutes a challenge. While this thesis comes a long way by systematically fusing data sources, future research, the research community and, more generally speaking, society, stand to benefit from improvements in data sharing. For this reason, this thesis also calls for structural improvements in data sharing.

Samenvatting

De maatschappij hangt tegenwoordig sterk af van het Internet voor globale communicatie. Hoewel kernonderdelen van het Internet ooit zijn ontworpen met weerstandsvermogen in gedachte worden de stabiliteit en duurzaamheid van het Internet in deze tijd voortdurend onderworpen aan opzettelijke bedreigingen. Onder deze bedreigingen vallen zogeheten Denial-of-Service (DoS) aanvallen, een type aanval met mogelijk zeer ellendige gevolgen.

Het algemene publiek begon circa tien jaar geleden (in 2010) een beeld te vormen wat de mogelijke gevolgen van DoS aanvallen inhouden. Dit was nadat een reeks aanvallen door WikiLeaks supporters voor ontwrichting op het Internet had gezorgd. Meerdere financiële organisaties werden destijds aangevallen en in veel gevallen werden Web sites offline geforceerd danwel merkbaar verstoord. Een paar jaar later, in 2013, shockeerde de *Spamhaus* aanval velen omdat het daarbij betrokken aanvalsvolume (van netwerk verkeer) record-brekend was. Nog enkele jaren verder (in 2016) zorgde de aanval op *Dyn*, een leverancier van DNS diensten, voor significante storing. Vele diensten die van *Dyn* afhankelijk waren werkten niet en dit was merkbaar in groten getale, voornamelijk onder gebruikers in Noord Amerika en Europa. De voorgenoemde voorbeelden zijn slechts een selectie van gepubliceerde incidenten die de ernst van de DoS dreiging benadrukken. Hoewel het DoS probleem niet nieuw is hebben we voornamelijk over de afgelopen jaren een sterke toename in het aantal alsmede de intensiteit van aanvallen waargenomen. De lat voor het behaalde aanvalsvolume wordt steeds hoger gelegd. Recente aanvallen hebben naar verluidt volumens van 1.7 Tbps behaald, waardoor de ooit shockerende *Spamhaus* aanval met 300 Gbps nu slechts kinderspel lijkt. Sommige experts menen ook dat we de volle potentie van aanvallen nog niet gezien hebben, wat tot de vraag leidt: hoeveel record-brekende aanvallen gaan er in de komende jaren nog berucht worden?

Jaarlijks verliezen ondernemingen honderden miljoenen euros als gevolg van aanvallen. Als het op vitale infrastructuur neerkomt dan staan de nationale veiligheid en mogelijk zelfs mensenlevens op het spel. De zich doorontwikkellende DoS dreiging maakt effectieve middelen voor bescherming (ofwel mitigatie) uiterst noodzakelijk. Nogmaals, hoewel het DoS probleem niet nieuw is, is de schaal van het probleem nog steeds onduidelijk. We weten niet precies waar we ons tegen verdedigen en om het probleem frontaal aan te kunnen pakken

is een beter begrip vormen een vereiste. Vele technische en maatschappelijke uitdagingen moeten worden opgelost ten behoeve van situationeel bewustzijn.

De wereld staat uiteraard niet stil. De opkomst van het DoS probleem heeft niet alleen voor de ontwikkeling van diverse mitigatie technieken gezorgd, maar ook tot een lucratieve markt voor commerciële producten geleid. Ondernemingen en andere potentiële gebruikers van beschermingsmiddelen worden geconfronteerd met verscheidene keuzes waarvan de best passende keuze niet altijd voor de hand ligt. Tevens kan het gebruik van zulke middelen verborgen nadelen met zich meebrengen die beter begrepen moeten worden.

Beleidsmakers en overheden staan ook voor vraagstukken. Wat moet er gebeuren om de nationale cyberveiligheid te verbeteren? Moeten de markt voor beschermingsmiddelen gestimuleerd worden? Brengt centralisatie rondom een paar aanbieders problemen met zich mee? En kunnen we als land (te) afhankelijk van andere landen worden als het gaat om het beschermen van vitale infrastructuur? Zulke vragen maken duidelijk dat het DoS probleem ook maatschappelijke problemen met zich meebrengt.

Omdat het vormen van een beter begrip van het DoS probleem vereist is om de algemene Internet veiligheid te verbeteren heeft dit proefschrift drie hoofdbijdragen. Ten eerste voert dit proefschrift een grondige, grootschalige karakterisatie van aanvallen en aanvalsdoelen uit om een beter inzicht te krijgen in waar we ons tegen verdedigen. Ten tweede verbetert dit proefschrift kennis over de Internet-brede ingebruikname van diverse beschermingsmiddelen alsmede de wijze waarop deze worden gebruikt. Ten derde onderzoekt dit proefschrift verborgen nadelen van beschermingsmiddelen die bij verkeerd gebruik de effectiviteit van mitigatie kunnen ondermijnen.

Met betrekking tot de eerste bijdrage onthult dit proefschrift de massieve schaal van het DoS probleem. We identificeren en fuseren op systematische wijze diverse data van onafhankelijke, globale Internet meetinfrastructuren om een macroscopische karakterisatie van aanvallen en aanvalsdoelen uit te voeren. We stuiten op bijna 21 miljoen aanvallen over een periode van twee jaar. We tonen ook aan dat, gedurende twee jaar, circa één derde van alle /24 netwerk adres blokken die naar schatting actief op het Internet worden gebruikt het doelwit van een DoS aanval zijn geweest. Dit proefschrift kijkt ook naar de mogelijk gevolgen van aanvallen. We laten zien dat Web hosting infrastructuur prominent wordt aangevallen en dat de aangevallen structuur collectief met meer dan 130 miljoen Web sites kan worden geassocieerd (gedurende twee jaar).

Voor de tweede bijdrage onderzoekt dit proefschrift twee beschermingsmiddelen: zogeheten cloud-gebaseerde diensten en BGP blackholing. We quantificeren de ingebruikname van cloud diensten op globale schaal en laten hierin een prominente trend zien. Onze resultaten tonen onder domeinnamen in de **com**, **net** en **org** zones een relatieve groei in ingebruikname aan van $1.24\times$ (gedu-

rence 1.5-jaar). Tezamen representeren deze zones circa de helft van alle globale domeinnamen. We onderzoeken ook tot in hoeverre slachtoffers van DoS aanvallen diensten ingebruiknemen na een aanval (we noemen dit migratie). Onze resultaten tonen aan dat de intensiteit van een aanval een belangrijke factor is voor migratie, terwijl herhaling en de duur van aanvallen dat niet zijn. Qua BGP blackholing onderzoekt dit proefschrift verscheidene operationele aspecten ‘in het wild’. Onze resultaten tonen aan hoe blackholing in de praktijk wordt ingezet door operatoren. Voor bijna 4% van DoS aanvallen die met blackholing werden gemitigeerd duurde het langer dan 24 uur nadat de aanval gestopt was eer operatoren het ingezette middel terugtrokken. Gedurende deze tijd zijn de beschermde machines niet bereikbaar voor (delen van) het Internet. Het ogenschijnlijke gebrek in automatisch herstel leidt tot zorgen dat machines alsmede diensten die op deze machines draaien te lang van gebruikers zijn afgesneden. We tonen daarbij ook aan dat minder intense aanvallen ook met blackholing worden gemitigeerd: in 13% van de gevallen was het (afgeleide) volume van de aanval slechts 3 *Mbps*. Gezien blackholing als ‘zelf-toegebrachte’ DoS kan worden beschouwd leiden deze resultaten tot de vraag hoe weinig moeite aanvallers moeten doen om operatoren dit drastische middel in te laten zetten.

Voor de derde bijdrage onderzoekt dit proefschrift nadelen in het gebruik van beschermingsmiddelen die de effectiviteit van mitigatie kunnen ondermijnen. Cloud-gebaseerde diensen kunnen door geraffineerde aanvallers worden omzijd als gevolg van fouten in gebruik. Niet alle gebruikers zien deze fouten mogelijk in, wat tot een vals gevoel van veiligheid kan leiden. We quantificeren dit nadeel op het Internet door op de meest populaire Web sites ter wereld en op vooraanstaande commerciële beschermingsdiensten te focuseren. Onze resultaten benadrukken de omvang van het probleem: de bescherming van 41% van de beschouwde Web sites kan worden omzijd. Wat betreft blackholing neemt dit proefschrift de eerste stappen om te onderzoeken tot in hoeverre machines onnodig van het Internet worden afgesneden. We quantificeren dit in termen van alledaagse Internet diensten die op de getroffen machines draaien.

Dit proefschrift zal van begin af aan laten zien dat een rudimentaire uitdaging betrekking heeft op data. Het vergaren en ontwikkelen van diverse (ruwe) data bronnen om vervolgens methodologisch binnen de DoS context te bestuderen zorgt voor een uitdaging. In dit proefschrift komen we een heel eind door systematisch data bronnen te fuseren om de hoofdbijdragen te bewerkstelligen. Dat laat niet weg dat toekomstig onderzoek, de wetenschappelijke gemeenschap, en breder genomen de maatschappij in het algemeen, voordelen kunnen halen uit het delen van data. Om deze reden roept dit proefschrift op tot structurele verbeteringen in het delen van data. En in het licht van deze oproep, alsmede om de basis te leggen voor onderzoek dat op dit proefschrift voortbouwt, hebben we een uitgebreide data set van DoS aanvallen publiek beschikbaar gemaakt.

Contents

1	Introduction	1
1.1	Distributed Denial-of-Service Attacks	1
1.2	Mitigating DDoS Attacks	2
1.3	Challenges with Mitigation Solutions	4
1.4	Goals, Research Questions and Approach	9
1.5	Organization and Key Contributions	13
2	Background on DoS Attacks and Mitigation	19
2.1	Reading Guide	19
2.2	(Distributed) Denial-of-Service Attacks	19
2.3	Attack Mitigation	22
2.4	BGP-Based Network Traffic Diversion	25
2.5	Moment of Mitigation	26
3	Attack Characterization	30
3.1	Introduction	30
3.2	Data Sources on DoS Activity	31
3.3	Data Sets	35
3.4	Analysis of Attacks	36
3.5	Related Work	46
3.6	Concluding Remarks	47
4	Impact of Attacks	48
4.1	Introduction	48
4.2	DNS Measurement Data	49
4.3	The Effect of Attacks on the Web	51
4.4	Related Work	56
4.5	Concluding Remarks	56
5	DDoS Protection Services	58
5.1	Introduction	58
5.2	Data Sources on DPS Use	59
5.3	Methodology and Data set	60

5.4	Adoption and Characteristics of Use	62
5.5	Attack Effects on Adoption	71
5.6	Concluding Remarks	78
6	BGP Blackholing	79
6.1	Introduction	79
6.2	Methodology and Data Sources	81
6.3	Blackholed Attacks	89
6.4	Blackholed Services	96
6.5	Blackholing Efficacy	99
6.6	Corroborated Collateral Damage	102
6.7	Related Work	104
6.8	Concluding Remarks	104
7	Exposure to Direct Attacks	106
7.1	Introduction	107
7.2	Data Sources	109
7.3	Exposure Vectors	109
7.4	Methodology	112
7.5	Data sets	116
7.6	Results	118
7.7	Related Work	124
7.8	Concluding Remarks	124
8	Conclusions	126
8.1	Main Conclusions	126
8.2	Research Questions Revisited	131
8.3	Prospects for Future Research	138
	Appendices	139
	Bibliography	139
	About the Author	149

Introduction

Our primary communications fabric is under siege. The evolution of the Internet has had a revolutionary impact on modern society. What started as a technology to interconnect educational institutes, research centers and alike has – over the past three decades or so – taken over global communications. The Internet has become an integral part of modern society, tying into, among others, commerce, technology and entertainment. We use the Internet to communicate with others through instant messaging, e-mail or voice over Internet calls. And we rely on it to both find and disseminate important information, for example by accessing news on-line, or by communicating with government. Due to the Internet's omnipresence, life as most of us know it is unthinkable without it. As we have a dependency on the Internet for communication, its availability – taken for granted by many – is of vital importance. Although critical components of the Internet were originally designed with resilience in mind, the stability and reliability of the Internet are nowadays continuously subject to deliberate threats, including devastating Denial-of-Service (DoS) attacks.

A rigorous characterization of the DoS phenomenon, and of countermeasures to mitigate the associated risks, is missing and faces many analytic challenges. This thesis addresses precisely this open issue, by taking a measurement-based approach to characterizing attacks and mitigation solutions. Our work advances situational awareness universally, and demonstrates our ability to inform Internet research, network operations and policy makers about the growing DoS threat.

1.1 Distributed Denial-of-Service Attacks

Over the past decades, DoS attacks have rapidly increased in terms of occurrence and intensity, steadily becoming one of the largest threats to the stability and reliability of the Internet. In this thesis we reveal the massive scale of the problem, by showing, among others, that one-third of all /24 networks estimated to be active on the Internet have suffered at least one DoS attack during a recent two-year observation period.

As strongly suggested by the name, DoS attacks are used by attackers to achieve denial of service. In essence, this entails cutting a networked service entirely off the network, e.g., the Internet, by any means possible. As an example, consider DoS attacks against on-line news media outlets or banks, scenarios that are not merely fictional but in fact have become reality in various notorious cases [26, 60, 86]. The motivations of attackers can vary wildly, including – but certainly not limited to – creating a distraction from other malicious activity (e.g., masking data theft [41, 66]), hacktivism (e.g., politically motivated attacks) [36, 58], or cyber-extortion (e.g., threatening a banks to take down its e-banking application unless a ransom is paid) [91].

Attacks can come in various shapes and sizes. In this thesis we present a large-scale characterization of attacks. For this introduction it is important to note that many types of attacks put not only a burden (the intended burden) on the target of the attack (i.e., the intended victim), but also threaten interconnecting network links. In case attacks are distributed, attack traffic will originate from multiple locations. Before converging on the target, this traffic may have adverse effects on globally disperse network segments. Moreover, for some types of attacks (including reflection attacks), core Internet infrastructure is abused to bring about the attack, which means that services that are essential for Internet operation may be involved in the attack even though they are not the intended target. As a consequence, DoS is not only a threat to the attack target and the interconnecting network infrastructure, but potentially also to core Internet services. Our large-scale characterization of attacks will underpin that attacks that abuse core Infrastructure, at times, are launched jointly with other attack types, savagely.

The collapse of any component involved in an attack may have ripple effects, create cascading failure, and potentially have an immense impact on the Internet [44]. In the face of the DoS threat that is nowadays an unwanted reality, effective defenses are an absolute necessity.

1.2 Mitigating DDoS Attacks

The upsurge of DoS attacks has given rise to the development of many diverse mitigation solutions. In this thesis we study two global solutions: cloud-based protection services and BGP blackholing.

There are types of solutions that operate close to the assets that they are meant to protect. For example consider an “in-line” appliance (e.g., firewall) that is placed in front of (and local to) a Web server that needs protection. Other types of solutions work in a distributed fashion, using load-balancing and network traffic diversion techniques, potentially on a global scale. We provide

more details on DoS mitigation solutions later, but stress here that generally speaking, on the one hand, defending against DoS attacks is better done closer to the Internet backbone, before malicious network traffic has a chance to do real harm. No strictly “in-line” solution is capable of thwarting attacks the largest of attacks in terms of network traffic volume. On the other hand, detection is generally better done closer to the target, where malicious traffic from potentially diverse origins converges and starts doing harm [78]. Because of this, various proven solutions (including the two studied in this thesis) are inter-domain, meaning that telemetry information for detection as well as reactive control measures for mitigation are exchanged across organizational boundaries. With some types of solutions protection is outsourced to other parties altogether, for example when a DDoS Protection Service (DPS) is contracted to offer a “cloud-based” solution.

Figure 1.1: Growth (relative) of cloud-based protection services use in .nl, over the period March 2016 – June 2018 (source: [88])

Choosing a suitable mitigation solution is a challenge in itself. When it comes to protecting citizens and vital infrastructure against cybercrime including DoS, governments have a clear stake in promoting cybersafety. This includes fostering a market for mitigation solutions. At the request of the Netherlands Ministry of Justice and Security, the CPB Netherlands Bureau for Economic Policy Analysis recently released a report on cybersecurity and economics [88], on which we were asked to collaborate. The report assesses risk of cybercrime, among others. It notably includes an analysis of the market for mitigation solutions available to companies in the Netherlands. The report stipulates an uptake in (leading) cloud-based mitigation solutions among Web sites with a Netherlands domain name (.nl). Figure 1.1 (this is Figure 18 in the CPB report) shows a relative growth in DPS use of $1.32\times$ among Dutch Web sites over 27 months. Most mitigation providers are US-based, which gives rise to concerns about digital dependence (more on this later). It is important to note that we performed the

underlying analysis for the CPB report. The market analysis was made possible by the results of this thesis.

Which mitigation solution fits best in essence varies on a case by case basis. The types of attacks that one may have to deal with – as well as the consequences of successful attacks – are considerations, but choosing a mitigation solution is not always easy. There are many other circumstances to consider. For example, a bank may not want to allow a third party to inspect confidential communication between the bank and its customers (e.g., e-banking activity). Yet some types of attacks can only be detected by inspecting unencrypted communication. As a consequence, the bank will need to detect some types of attacks in their own data center, where encrypted connections terminate. On the other hand, the same bank may have to deal with sizable volumetric attacks that cannot be dealt with merely in their data center. All things considered, the bank may opt to go with a hybrid solution.

An organization that has fewer concerns relating to confidentiality may fully rely on a cloud-based solution and hand over the keys to decrypt network traffic. As a final example, an organization that accepts the risk of downtime following attacks that would be expensive to mitigate through outsourcing may choose to only deploy an in-line mitigation appliance.

1.3 Challenges with Mitigation Solutions

There are many challenges when it comes to DoS mitigation, including but not limited to: (i) challenges in knowing exactly what it is we are defending against; and (ii) challenges in the adoption and operation of mitigation solutions. We contribute towards overcoming these challenges in this thesis. We successfully fuse data from diverse sources (e.g., attack telemetry) in pursuit of situational awareness. And armed with enriched, large-scale data, we reveal, among others, common mistakes in the deployment and operation of various global mitigation solutions.

1.3.1 Reporting at Scale

If we are to believe commercial providers of mitigation solutions, the scale of the DoS problem is immense. Many leading providers publish yearly or quarterly reports on attacks and attack trends. Imperva, for example, release a quarterly *Global DDoS Threat Landscape Report*. Their *Q4 2017* report contains a statistical analysis of 5000 network and application-layer attacks observed by their own infrastructure [23]. The report reports a near-doubling of application-layer attacks, a decline in network-layer attacks, and also reveals that the cryptocurrency industry (e.g., Bitcoin) had risen in the most-targeted ranking.

Akamai frequently releases *State of the Internet / Security* reports. Their *Q4 2017* release, for example, reports on a 14% increase in total DDoS attacks (compared to Q4 2016) and a 14% increase in network-layer attacks [77].

Reporting on attack characteristics *at scale* constitutes a challenge in which data availability and processing capability play significant roles. Our DoS attacks characterization in this thesis accounts for nearly 21 million attacks.

1.3.2 Availability and Integration of Diverse Data

It is important to note that reports such as those mentioned above are based on *data*, but those data are not only proprietary, but also specific to the customer bases of the providers in question. The methods used are often not included or not sufficiently explained. Moreover, it stands to reason that vendors of mitigation solutions stand to benefit from making the problem appear larger than it is. That is not to say that there are no academic works on quantifying the DoS problem. Many works, however, are outdated or limited in scope, focusing for example only on one category of attacks (e.g., reflection attacks), or on attack activity, albeit diverse, learned only from specific malicious infrastructure segments (i.e., botnet families). It is a challenge to identify and fuse these data to get a global view of the DoS phenomenon. In this thesis we address this by considering diverse and independent data sources that provide Internet-wide indicators of DoS activity, using open and established methodologies, where available. By successfully fusing these data we unveil eye-opening statistics about global attack activity, and demonstrate our capacity to inform network operators and policy makers. Additionally, to address the sparse availability of data, we make available to the research community attack data to stimulate further research on the DoS phenomenon.

1.3.3 Adoption of Mitigation and Expertise of Users

Even though diverse solutions are readily available to mitigate attacks, quantitative knowledge of their adoption on the Internet is limited. In addition, an understanding of how solutions are deployed and operated when operators are faced with attacks featuring differing characteristics is missing. A related challenge stems from the potential disconnect between the ease of setup of mitigation solutions and the expertise of adopting operators. Providers, be it of cloud-based services or on-site appliances, stand to benefit from a low adoption barrier. Often they try to capitalize rapid product (or service) deployment, because that is what companies need in times of crisis (i.e., when attacked). But what exactly does a black box with proprietary algorithms do after it is so easily plugged into a network? While that box may effectively mitigate attacks and

tempt its delighted new owners to leave it untouched, does not turning on some of the little knobs have any adverse effects down the line? Are there operational pitfalls that a small-to-medium enterprise (SME) without seasoned network operators and security engineers face when using certain mitigation technologies? In this thesis we highlight that mistakes indeed are made in the deployment and operation of mitigation solutions, which arguably leave some operators and users with a false sense of security.

Attackers may also try to seize on bad operational practices by users of mitigation solutions as an opportunity. Our work corroborates this notion by showing that cloud-based providers are, at times, bypassed by attackers.

On top of the challenges described thus far, there are *other*, societal challenges when it comes to the DoS problem. These challenges include, but are not limited to: *(i)* encouraging the development of cybersafety on a national level; and *(ii)* independent control over protection; and *(iii)* protecting the data privacy of citizens.

1.3.4 National Cybersafety

Again, to promote cybersafety on a national level, governments may want to foster the market for – and availability of – mitigation solutions. Equally important is informing citizens, companies and alike – not only to raise awareness about the actual size of the DoS threat, but also about possible solutions. If we are to believe the media, the DoS problem is significant. However, typically only record-setting attacks make the news, or attacks with high-profile targets. Is a SME as likely to get hit as a Fortune 500 company? Are companies that operate in sectors that are less attractive for, e.g., cyber-extortion, as likely to see their Web site get hit as banks? Companies may ask these questions before designating capital and operational expenditure to proactively adopt a mitigation solution.

1.3.5 Centralization and Digital Independence

Other challenges surround centralization and dependence on foreign providers. With a few large players dominating the market, the safety and security of a country may become dependent on foreign entities, for example when the means of a government to communicate with its citizens factors into the equation. The CPB report on cybersecurity and economics (mentioned previously in Section 1.2) reports on concentration around large providers. The report raises concern among policy makers, using Figure 1.2, that a majority (98.48%) of Dutch Web site operators tend to stick to a single mitigation provider over prolonged periods. The report stipulates that lack of diversification introduces

the risk of becoming dependent on foreign entities when it comes to national cybersafety matters. This part of the CPB report would (also) not have been possible without the work in this thesis.

Figure 1.2: Growth (relative) of cloud-based protection services use in .nl, over the period March 2015 – June 2018

A recent news article in the daily newspaper *Het Financieele Dagblad* echoes similar concerns about foreign dependence, especially when it concerns Dutch banks [49]. As it turns out, various banks in the Netherlands – and about half of the world’s largest banks for that matter – depend on *one specific* mitigation provider. The work in this thesis contributed to the article.¹

1.3.6 Privacy of Citizen Data

Finally, outsourcing protection to foreign providers also brings about concerns about confidentiality and privacy. A commercial provider may be subject to various territorial jurisdictions if it operates (network infrastructure) in multiple countries. And its customers may have limited to no control over where traffic is routed. This means that the diversion of customer traffic may subject it to inspection by foreign parties (e.g., intelligence agencies). To make matters worse, a provider may be subpoenaed (in secrecy) to share data (e.g., through a FISA warrant in case the USA Freedom Act applies), either because it is registered in a foreign country (the US for example), or merely because it also does business there. Even in cases where third parties see only encrypted traffic, a great deal of information can still be learned. For example, the origin and even identity of (legitimate) clients of, e.g., a bank can be determined based on lower-layer network connection properties. As another example,

¹Please note that our contributions are of a statistical form and do not extend to any editorialisations in the article.

behavioral patterns can be learned based on connection timing and size.

1.3.7 The Battles We Pick

These challenges cannot all be solved at once. We argue that many societal challenges cannot be tackled without advancing our technical understanding about the DoS phenomenon first. In this thesis we focus on various technical challenges. We address analytic challenges that relate to data scale and availability and processing. We use the resulting data to scientifically research and characterize parts of the DoS ecosystem. We advance knowledge about the adoption of mitigation solutions. And we further the understanding of operational use of mitigation solutions.

1.4 Goals, Research Questions and Approach

In Section 1.3 we overviewed problems surrounding scientific reporting on the scale and characteristics of the DoS problem. On the basis of these problems, we define the first research goal of this thesis as follows:

***Goal 1:** to study the DoS phenomenon on a global, Internet-wide scale, and to identify, join and validate – where applicable – existing data to broadly report on the the DoS problem*

We also pointed out that there is limited knowledge about mitigation solutions within the research community. The missing puzzle pieces include an understanding of adoption at scale, as well as an understanding of how solutions operate in the face of attacks. For this reason we define the second research goal of this thesis as follows:

***Goal 2:** to study the adoption of DDoS Protection Services and BGP black-holing, and to investigate the operational practices of operators that use these solutions*

Finally, we argued that lack of expertise on the part of users of mitigation solutions may lead to mistakes in deployment and operation. These can lead to undesirable side effects, create a false sense of security, and may even be seized on by attackers as an opportunity. These notions lead us to the third and final research goal of this thesis:

***Goal 3:** to study problems surrounding the use of mitigation solutions that result from mistakes in use and bad operational practices, and to investigate whether or not attackers seize on these as an opportunity*

In the three sections that follow we break our three goals down into research questions. We also summarize our approaches to answering each and every one of the research questions.

1.4.1 Goal 1: The DoS Phenomenon

Research Questions

In the first goal we expressed wanting to study the DoS phenomenon on a global, Internet-wide scale. This leads to our first research question:

***RQ 1:** Which data sources on DoS do we need in order to study the DoS phenomenon on a global scale? Are there existing data that we can work with, fuse or derive from? Or do we need to gather new measurements?*

We address RQ 1 in nearly every chapter of this thesis – mainly Chapters 3 through 7 – as we incrementally add data before using it to expand our study of the DoS ecosystem.

Once we have identified data, the next thing to ask is what the DoS ecosystem looks like on a global scale in terms of attacks. Do attacks occur as often as commercial reports suggest? And which attack types are common? This leads to our second research question:

***RQ 2:** What does the DoS landscape look like on a global scale in terms of attack occurrence and attack types?*

We address RQ 2 in Chapter 3 of this thesis.

Once we have identified the scale of the DoS problem in terms of attack occurrence and attack types, the question that naturally follows relates to the attacked targets. As such, our third research question is:

***RQ 3:** Which targets are involved in DoS attacks? And what is the potential impact that attacks have on these targets?*

We address RQ 3 in Chapters 3 and 4.

Approach

To address the research questions above, we take a measurement-based approach. It stands to reason that it is impossible to study the DoS problem at an Internet-wide scale based on anything but global Internet measurement data. We use large-scale passive and active measurements from diverse vantage points all over the world, to gather a variety of independent data types. Given the challenge of processing such data, we will fuse, derive, and analyze data sets by applying Big Data Analytics. In the process, we will identify and verify, where applicable, pre-existing methodologies to measure, as well as devise new ones along the way where necessary. To enable reproducibility and future research, and to defeat the limitations of some of the existing reports on DoS activity, we will make available data to other researchers.

1.4.2 Goal 2: DoS Mitigation Solutions

Research Questions

The second goal of this thesis is to study the use of diverse mitigation solutions, which includes their adoption as well as factors that drive their use. We define the first research question towards meeting this goal as:

***RQ 4:** Can we quantify the adoption of commercial, cloud-based DDoS Protection Services? In which manner do customers use such services? And what are the factors that drive adoption?*

We address RQ 4 in Chapter 5 of this thesis.

As the focus of this thesis is on diverse mitigation solutions that we can empirically study on an Internet-wide scale, our next research question is:

***RQ 5:** How widespread is the use of BGP Blackholing for the purposes of DoS mitigation? And how do users, i.e., network operators, use blackholing when faced with DoS attacks?*

We address RQ 5 in Chapter 6.

Approach

As is the case for the first goal, our approach to answering the research questions for the second goal is systematic, large-scale and measurement-based. We want to study mitigation solutions that deal with the DoS phenomenon on a global scale. Consequentially, our focus will be on proven, Internet-wide strategies that cross organizational boundaries.²

With a shift from attacks to mitigation, we have to identify and add new data to study mitigation solutions. We will fuse these data with the previously identified data on attacks to come to insights on mitigation practices following attacks, among others.

1.4.3 Goal 3: Problems with Mitigation

The third and final goal of this thesis is to study potential problems with mitigation. We previously argued that a lack of experience in properly setting up a mitigation solution, as well as bad operational practices, may have undesirable side effects. Our sixth research question is therefore:

***RQ 6:** Can we identify problems with the adoption of DDoS Protection Services? Can we quantify this problem on the Internet? And do we see evidence that attackers actively seize on potential problems?*

RQ 6 is addressed in Chapter 7.

Blackholing, by design, is a rather coarse-grained approach to attack mitigation. It is effectively an intentional “self-inflicted” Denial-of-Service. As we want to study the problems with blackholing as well as DDoS Protection Services, our seventh and final research question is:

²This scope excludes strategies that are intra-domain (e.g., strictly in-line appliances).

***RQ 7:** Can we quantify the adverse effects of blackholing on the Internet?*

RQ 7 is addressed in Chapter 6.

Approach

Our approach to answering the final research questions is in line with the previous approaches. We focus on the same diverse mitigation solutions as before and will identify, fuse and analyze new data, as well as devise methodologies, to reach the third goal. We developed in this process active measurement infrastructures to gather specific insights about blackholed prefixes and users of protection services.

1.5 Organization and Key Contributions

Figure 1.3 shows a schematic of the structure of this thesis. The schematic shows the relation between chapters, as well as how chapters compose distinct parts of this thesis. The relations between chapters suggest a reading order, which means, for example, that readers are recommended to read Chapter 3 before reading Chapters 5 and 6. In the sections that follow, for each chapter, we provide a summary, list key contributions, and refer to the publications on which the chapter is based.

Figure 1.3: Thesis structure schematic

Chapter 2: Background on DoS Attacks and Mitigation

In this chapter we provide background information on (Distributed) Denial-of-Service (DDoS) attacks and on DDoS attack mitigation. We start with a brief history of the rise of DDoS attacks. We outline various categories of attacks as well as commonly used attack types. We discuss various solutions for attack mitigation. Finally, we provide technical background information for the mitigation technologies considered in this thesis as an aid to help understand our measurement methodologies. This chapter in part is based on the following peer-reviewed publication, as well as on background sections of peer-reviewed publications that are referenced under later chapters:

- M. Jonker and A. Sperotto. *Mitigating DDoS Attacks using OpenFlow-based Software Defined Networking*. In Proceedings of the 9th IFIP WG 6.6 International Conference on Autonomous Infrastructure, Management, and Security (AIMS'15). Ghent, Belgium [55].

Chapter 3: Attack Characterization

This chapter discusses our first steps towards addressing challenges that pertain to: (i) data availability; and (ii) processing large-scale and diverse data

sources. In addition, we advance through a rigorous characterization of attacks to understand what we (collectively speaking) are defending against. The main contributions of this chapter are that we:

- Establish a novel approach that enables a more thorough scientific approach to characterizing the Denial-of-Service ecosystem;
- Use our approach to systematically fuse diverse data from independent, global Internet measurement infrastructures;
- Perform the first macroscopic characterization of both DoS attacks and attack targets at scale;
- Demonstrate the potential of our approach to provide situational awareness and inform Internet research, network operation and policy communities about a growing threat to Internet stability and reliability.

The results of the work discussed in this chapter were presented at an annual workshop that promotes discussion between academics, industry, and policy-makers on active Internet measurement. The goals of our presentation were to: *(i)* disseminate that the scale of the DoS problem is larger than previously reported; and *(ii)* report on experiences and the potential of fusing measurement infrastructures [32] (AIMS 2017). The framework we established paved the way for new research on DoS attacks and Internet security, even multi-disciplinary. An example of such is a study in collaboration with political scientists on the use of DoS attacks as a tool in non-democratic regions [72]. In another example our results have laid the groundwork for new research into DNS security and stability [9]. Finally, to facilitate access to independent researchers, as well as to make possible reproducibility, we published our data set through *IMPACT* [3]. *IMPACT*, short for *Information Marketplace for Policy and Analysis of Cyber-risk and Trust*, is a platform that “supports the global cyber-risk research & development community by coordinating and developing real-world data and information-sharing capabilities between academia, industry and government.” [7]

This chapter is based on the following peer-reviewed publication:

- M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto and A. Dainotti *Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem*. In proceedings of the 2017 ACM Internet Measurement Conference (IMC’17). London, United Kingdom [51].

The publication on which this chapter is based has received recognition in the following forms:

- The paper was among a handful selected nominations for the Dutch Cyber Security best Research Paper award (DCSRP2018 @ICT.OPEN2018);
- Some of the publication’s findings were covered by more than ten US and NL media outlets, which notably includes *TheRegister* and *Tweakers* [31, 109].

Chapter 4: Impact of Attacks

In Chapter 3 we established a novel framework to characterize the DoS ecosystem at scale. Our analysis of DoS attacks showed that among attacked targets, Web infrastructure is prominent. In this chapter we evaluate the impact potential of DoS attacks on the Internet, focusing on the Web. Furthermore, we study the potential for Web sites to become collateral damage of a DoS attack by being co-hosted on shared infrastructure. With respect to the previous chapter, the main contributions of this chapter are that we:

- Illustrate the potential impact of DoS attacks by fusing an additional data source (i.e., active DNS measurements) in our framework;
- Unveil that Web infrastructure that belongs to large hosters is prominent among the attacked targets, and that targets sometimes involve millions of co-hosted Web sites;
- Show that for Web infrastructure targets, attackers are more likely to target protocols and ports specific to Web services;
- Reveal that over an extended period, about two-thirds of all Web sites found under the largest top-level domains can be associated with attacked hosts.

This chapter is based on (part of) the following peer-reviewed publication:

- M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto and A. Dainotti. *Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem*. In proceedings of the 2017 ACM Internet Measurement Conference (IMC’17). London, United Kingdom [51].

Chapter 5: DDoS Protection Services

This chapter focuses on DDoS Protection Services. The use of a DPS is the first among two global mitigation strategies discussed in this thesis. We study the adoption of protection services on the Internet, by inferring DPS use amongst a

representative and significant number of domain names. We also jointly analyze our new data source on DPS use with attacks data to shed light on factors influencing the adoption of protection services following attacks (we refer to this as “migration”). The main contributions of this chapter are that we:

- Quantify the use of protection services among more than 50% of all domain names in existence, for the largest commercial providers, revealing a prominent trend in adoption;
- Reveal that large parties such as Web hosters drive adoption and may dynamically divert network traffic for many Web sites at once, making potentially impactful decisions on behalf of the customer;
- Quantify the extent to which Web sites migrate after having been targeted by a DoS attack. We reveal that Web sites for which we observe an attack are more likely to migrate than those for which we do not, and show that repeated attacks and attack duration were non-determinative factors for migration, whilst a higher attack intensity was;
- Validate diverse methodologies that measure DoS attacks. First, by connecting, through data fusion, inferred attack activity to migration. And second, by validating the correctness of inferred attack attributes.

The results of the work discussed in this chapter were used as input for a risk report on cybersecurity and economics, written by the CPB Netherlands Bureau for Economic Policy Analysis, to inform – and at the request of – the Netherlands Ministry of Justice and Security [88]. Part of the work also contributed to a news article in the daily newspaper *Het Financieele Dagblad* that stipulates that the e-banking communication between citizens of the Netherlands and their banks may be trivially accessed by foreign entities.

This chapter is based on the following peer-reviewed publications:

- M. Jonker, A. Sperotto, R. van Rijswijk-Dei, R. Sadre and A. Pras. *Measuring the Adoption of DDoS Protection Services*. In proceedings of the 2016 ACM Internet Measurement Conference (IMC’16). Santa Monica, California, USA [57];
- M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto and A. Dainotti. *Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem*. In proceedings of the 2017 ACM Internet Measurement Conference (IMC’17). London, United Kingdom [51].

Chapter 6: BGP Blackholing

This chapter focuses on BGP blackholing, which is the second global attack mitigation strategy that we study in this thesis. We study the use of this countermeasure following DoS activity by fusing BGP data in our framework, and in doing so are able to shed light on operational aspects of mitigation. The main contributions of this chapter are that we:

- Identify and reveal operational aspects of BGP blackholing at scale, with several revelations that raise concern that hosts may be unnecessarily cutoff from the Internet by operators;
- Further validate preexisting methodologies through fusing and analyzing diverse data sources by, among others: (i) linking inferred attack activity to blackholing; (ii) linking blackholing to inferred filtering of network traffic; and (iii) validating (further) the correctness of inferred attack attributes;
- Quantify the extent to which blackholing may cutoff the common Internet services Web, mail and DNS (we refer to this as “service collateral damage”) and present and apply a methodology based on reactive measurements to corroborate collateral in specific cases.

In addition to a presentation at the main publication venue, the results of the work discussed in this chapter were disseminated to research and operator communities in various other forms. First, based on preliminary findings, awareness of service collateral damage was raised at an annual workshop that promotes discussion between academics, industry, and policymakers on active Internet measurement [33] (AIMS 2018). Second, results were presented at RIPE78, a networking conference where network operators, Internet service providers and alike could be informed about the bad operational practices and the quantified drawback of blackholing [12]. Third, awareness was raised through a blogpost at APNIC, also targeting network operators and alike [53]. The work also enabled collaborative research on the potential of a less coarse-grained, but not yet widely adopted form of mitigation, i.e., BGP FlowSpec. This work became runner up in the ACM SIGCOMM Student Research Competition (SRC) 2018 [45].

This chapter is based on the following peer-reviewed publication:

- M. Jonker, A. Pras, A. Dainotti and A. Sperotto. *A First Joint Look at DoS Attacks and BGP Blackholing in the Wild*. In proceedings of the 2018 ACM Internet Measurement Conference (IMC’18). Boston, Massachusetts, USA [52].

Chapter 7: Exposure to Direct Attacks

In Chapter 6 we studied an inherent drawback of using BGP blackholing for DDoS mitigation: services becoming collateral damage. In this chapter we study a major drawback of using DDoS Protection Services: their common bypassability as a result of so-called “origin exposure”. Origin exposure involves supposedly hidden DPS customer infrastructure (i.e., IP addresses) becoming known to outsiders. The main contributions of this chapter are that we:

- Identify a comprehensive set of vectors through which origins of DPS customers can be exposed, including novel vectors not previously reported in literature, and use this set to quantify origin exposure at scale, for the world’s most popular Web sites, and for leading commercial DDoS Protection services;
- Unveil the scale of the bypassability problem: *41%* of *11 k* Web sites considered exposed their origin through at least one vector;
- We match vulnerable DPS customers with data on DoS activity, providing for the first time a look at whether attacks actively bypass protection, and showing high-intensity attacks on *19%* of exposed Web sites.

Early results of the work discussed in this chapter were discussed at various workshops. First, the ongoing research effort was discussed at the *3TU Cyber Security Workshop 2016*, allowing for feedback from peers to help steer the work. Second, the work was discussed at the *DNS and Internet Naming Research Directions (DINR) 2016* workshop on challenges in the DNS (attended by academics and network operators [22]), both to raise awareness, as well as to allow for feedback to steer further investigation.

This chapter is based on the following peer-reviewed publication:

- M. Jonker and A. Sperotto. *Measuring Exposure in DDoS Protection Services*. In proceedings of the 13th International Conference on Network and Service Management (CNSM’17). Tokyo, Japan [56].

Chapter 8: Conclusions

Taking the research discussed in all of the previous chapters into account, we draw conclusions in the final chapter of this thesis. In addition, we outline possible directions for future research.

Background on DoS Attacks and Mitigation

The purpose of this chapter is to give the reader basic background information on various concepts within the (Distributed) Denial-of-Service (DDoS) attack landscape. Specifically, we provide an introduction to attacks and attack mitigation.

2.1 Reading Guide

This chapter is intended to serve the reader basic background information on Denial-of-Service attacks. We will start by providing a brief history on the rise of the DDoS problem. Then we will outline various categories of attacks and attack types. Afterwards, we will address attack mitigation. We will focus on DDoS Protection Services (DPS) and BGP blackholing in particular as these are the two global mitigation solutions that we study in this thesis.

While the predominantly measurement-based work in this thesis uses a range of diverse data sources and, at times, established methodologies, we provide background information on these concepts in the chapters of first use, rather than in this background chapter.

2.2 (Distributed) Denial-of-Service Attacks

Denial-of-Service attacks, which have rapidly increased in frequency and intensity, are known to be used against anything ranging from home network devices

to core Internet infrastructure. DoS attacks can abuse core parts of the Internet infrastructure (e.g., the Domain Name System (DNS)). In some cases attacks do not require an underlying botnet. Ever-increasing records [81, 90] underpin that DoS attacks have become a significant threat to Internet reliability and stability. While sub *Tbps* attack network traffic volumes were considered shocking by many only a few years ago, such figures can nowadays be considered rear-view mirror limits when it comes to high-profile cases.

2.2.1 Years of Escalation

Denial-of-Service attacks have been long noted in the literature, but it was not until a large group of attacks referred to as “*operation payback*” by WikiLeaks supporters that the general public better understood the power of DoS attacks. As part of this wave of attacks in 2010, the Web sites of *MasterCard* and *Visa* were brought down entirely, and *PayPal*’s Web site was notably disrupted [26, 43]. Ever since, we have seen a rapid increase in DDoS attacks in occurrence and magnitude. The “*Spamhaus attack*” is a notorious example [74]. While its *300 Gbps* traffic peak created the largest-ever-seen DDoS attack at the time, it has since often been surpassed by more powerful attacks. Recent attacks have reportedly hit sheer attack traffic volumes of *1.7 Tbps* [81].

To make matters worse, the ability to launch attacks is nowadays no longer limited to people with advanced technical skills. The rise of the DoS-as-a-Service phenomenon (i.e., Booters) [59, 99], has dramatically expanded the population of potential perpetrators, who can now purchase, in exchange for mere pocket change, the execution of attacks powerful enough to saturate *1-10 Gbps* links.

Events such as the attack against Dyn [44] and a DNS root server [82] have demonstrated the vulnerability of critical Internet infrastructure to DoS attacks. The full potential of attacks has arguably yet to be seen. Leverett et al. [69] estimate the upper bound of distributed reflection and amplification attacks to be above *100 Tbps*, which prompts the question: how many record-breaking attacks have yet to reach notoriety?

2.2.2 Attack Types

Volumetric attacks

Attackers aim to disrupt services when they employ Denial-of-Service attacks, thereby causing harm to the service operator and legitimate users. DoS is commonly achieved through resource exhaustion, which can take place at the network level (e.g., by saturating a network link with packets) or at the server level (e.g., by overloading a networked daemon with requests). Such attacks are referred to as *volumetric* attacks as they involve a sheer mass of requests

to try to overwhelm a service. Depending on how attack traffic is generated, volumetric attacks can be divided into *direct* and *reflection* attacks.

Direct attacks

Direct attacks involve attack traffic sent directly to the target, originating from infrastructure under the control of the attacker. For example, an attacker can use his own machine, a compromised server, or a set of compromised devices (e.g., a botnet) under his or her command. To conceal infrastructure, to impede countermeasures, and to complicate attribution, attackers oftentimes employ random source IP address spoofing, i.e., setting source addresses in packet IP headers to a forged value.

Reflection attacks

In a reflection attack, third-party infrastructure (i.e., one or more reflectors) is abused to reflect attack traffic towards the victim. Reflection also involves source IP address spoofing, but it does not involve random address values. Rather, the attacker sets the source IP address of a request specifically to the victim's IP address. The reflector, which has no means of checking whether a request was sent legitimately or with a spoofed IP address when a connection-less protocol is used, then sends its response to the victim. Many protocols that allow for reflection also send responses that are much larger than the requests, causing the amount of reflected traffic sent towards the victim to be many times greater than that sent towards the reflector initially, i.e., it is amplified [98]. Amplification does not just affect older protocols such as NTP and IGMP [34, 100], but also newer protocols such as DNSSEC [108].

Semantic attacks

Next to volumetric attacks there are also *semantic* attacks. Semantic attacks do not necessarily aim for resource exhaustion but rather try to exploit flaws to deny access to a service. As an example consider the sending of a malformed request that crashes a networked daemon. This type of attack is tailored to work against a specific service, whereas volumetric attacks are mostly service-agnostic.

Volumetric attacks are nearly impossible to mitigate with strictly on-premise solutions because they operate at the network and transport layers [120]. This does not apply to semantic attacks, which have negligible bandwidth effects.

2.3 Attack Mitigation

In the face of the DDoS threat, diverse mitigation solutions have been – and continue to be – developed, including: dedicated appliances (e.g., those offered by Netscout Arbor [2] and Radware [4]); BGP Flowspec [6]; BGP blackholing [10]; and cloud-based services that can be contracted to “scrub” network traffic in their data centers (e.g., CloudFlare). The costs of adopting mitigation solutions varies, e.g., from the purchase and operating costs of an in-line appliance, to a flat-rate or pay-per-volume agreement with a cloud-based protection service.

Attacks can be mitigated on-site by means of a dedicated appliance [92], absorbed in the cloud, or be mitigated through a hybrid setup in which customer premises equipment is combined with a cloud-based component. Semantic attacks can be mitigated in-line [25, 78] whereas large volumetric attacks, i.e., very high in network traffic volume, are best mitigated closer to the Internet backbone. The reason for this is that attack traffic converges near the target, which is where the risk of congested network links is the highest. For this reason the detection of attacks is typically easier near or at the target [78]. For these reasons, various proven mitigation solutions are inter-domain.

The type of attack (i.e., volumetric or semantic) and the type of customer determine the potential of each mitigation approach. For example, banks may want to terminate encrypted e-banking connections themselves, and therefore require a hybrid solution in which an in-line appliance can decrypt connections and mitigate semantic attacks, while the cloud thwarts large volumetric attacks without being able to inspect confidential payloads.

In this thesis we focus on volumetric attacks and on two global, i.e., inter-domain mitigation solutions. The remainder of this chapter provides background information on the two mitigation solutions studied in this thesis.

2.3.1 DDoS Protection Services

The rise of DoS attacks has stimulated a market for DDoS Protection Services (DPSs), i.e., commercial parties that can be contracted to filter and drop malicious traffic before it reaches the intended target. Protection services thus offer victims of attacks a means to outsource protection to a knowledgeable party. A DPS can offer various types of mitigation solutions, meaning that they can deal with volumetric attacks, semantic attacks, or both. Solutions may require all network traffic to be diverted to the cloud (i.e., the security infrastructure of the DPS), or be based on a hybrid setup that also requires an in-line appliance at the customer. Moreover, the protection of a specific application or service (e.g., a Web site) can be outsourced, as well as the protection of entire networks.

Network Traffic Diversion

The key mechanism to outsource protection to a DPS is *network traffic diversion*, i.e., routing network traffic towards the security infrastructure of the DPS. One way to divert traffic to services that are reached on the basis of a domain name is to leverage the DNS, in a manner similar to how content delivery networks implement load balancing [48, 85]. An alternative is to use the Border Gateway Protocol to divert traffic towards the DPS infrastructure, in which the DPS announces a customer-used prefix to attract traffic.

In the next two sections we outline the functioning of widely used diversion mechanisms based on DNS and BGP, and how those are implemented in a DPS.

Figure 2.1: Schematic of DNS-based network traffic diversion

2.3.2 DNS-based Network Traffic Diversion

The DNS can be leveraged to divert network traffic. A prerequisite is that the protected host is reached on the basis of a domain name (e.g., a Web site). A DNS-based setup is typically combined with a reverse proxy. The reverse proxy is placed between potential sources of malicious attack traffic and the protected host. It is positioned to forward only benign traffic to the so-called *origin* – and also to serve responses on behalf of the origin, which need not be in the DPS infrastructure.

Figure 2.1 shows a DNS-based diversion setup. In this case the Web site `www.examp.le` is protected by the DPS. The domain name `www.examp.le` is configured to resolve to `10.0.0.1`, which is the IP address of the reverse proxy. Note that the proxy is located within the DPS infrastructure whereas the origin Web server – with IP address `172.16.0.1` – need not be. Any client looking to make a Web request (e.g., `GET`) should speak to `10.0.0.1`. Only `10.0.0.1` needs to be allowed to speak to the origin, which can be enforced with a properly configured firewall.

There are various ways in which the DNS can be leveraged to setup network traffic diversion. That is, using the example in Figure 2.1, there are various ways to direct clients to the reverse proxy at `10.0.0.1`. We will explain these methods next.

IP Address Only

The owner of a domain name can point the DNS to the reverse proxy. This comes down to setting an `A` record to point to a DPS-assigned IP address (i.e., that of the proxy). Depending on how the protection service operates, the IP address can either be customer-specific, or shared by multiple customers. Moreover, in case IPv6 is supported, an `AAAA` record can be set accordingly. An example of this method is shown below.

```
;; ANSWER SECTION:
www.examp.le      IN      A      10.0.0.1
;; AUTHORITY SECTION:
www.examp.le      IN      NS     ns.registr.ar
```

Note that `ns.registr.ar` is authoritative for `www.examp.le`, as indicated by the `NS` record. Also, the owner of `www.examp.le` has control over the IP address provided by `ns.registr.ar`.

Canonical Name

A second method uses a canonical name. `www.examp.le` can be made into an alias for another domain name by configuring a `CNAME` record. If the `CNAME` record of `x` references the canonical name `y`, then some record types of `x` are determined by the DNS zone of `y`. This means that the DNS operator of `y` can, among others, set `x`'s `A` records. In the example shown below, the domain `foob.ar` belongs to the DPS – and through so-called `CNAME` “expansion” controls the records of `www.examp.le`.

```
;; ANSWER SECTION:
www.examp.le      IN  CNAME   foob.ar
foob.ar           IN  A        10.0.0.1
;; AUTHORITY SECTION:
foob.ar           IN  NS       ns.foob.ar
```

The difference with the first method is that the owner of `www.examp.le` no longer directly configures the IP address records and thus did not set the IP address to `10.0.0.1`.

Delegation

The third method involves delegation. A owner of a domain can delegate its zone to the name servers of a DPS. In the example below, provided that the protection service operates `ns.foob.ar`, it is the DPS that is authoritative for `www.examp.le`. This means the DPS can set the A record to `10.0.0.1` – and controls any other record for that matter.

```
;; ANSWER SECTION:
www.examp.le      IN  A        10.0.0.1
;; AUTHORITY SECTION:
www.examp.le      IN  NS       ns.foob.ar
```

The subtle difference between the CNAME method and this method is that using the prior, a domain name owner does not have to altogether change the delegation.

When DNS-based network traffic diversion is used it is recommended to drop requests to the origin from anywhere except the reverse proxy – or a set of proxies [1]. Setting up a firewall for this purpose (see Figure 2.1) is not strictly necessary for operation and can therefore at times be neglected. Moreover, in some cases, setting up a firewall is a complicated or even infeasible endeavor if a large number of reverse proxies are used by a DPS [70]. The consequences are that the origin may be reached by attackers directly, which is something that we investigate in Chapter 7 of this thesis.

2.4 BGP-Based Network Traffic Diversion

The Border Gateway Protocol (BGP) can also be used to divert network traffic. In a BGP-based setup, the DPS announces a customer-used prefix, e.g., a /24, to divert all customer-destined traffic. All traffic destined for this customer net-

work is then routed towards the DPS infrastructure, where it can be analyzed and scrubbed. Clean traffic is then sent back to the customer by means of, e.g., a Generic Routing Encapsulation (GRE) tunnel. A BGP-based approach is typically used to protect entire networks or when a reverse proxy is not feasible. Figure 2.2 shows an example. The DPS announces `172.16.0.0/24`, the customer-used prefix is under protection. This ensures that all traffic is routed towards the security infrastructure of the DPS. After scrubbing, clean traffic is sent to the customer in a GRE tunnel.

Figure 2.2: Schematic of BGP-based network traffic diversion

Which method? – Customer needs weigh in on the potential of either approach to network traffic diversion. A hosting company that needs to protect their entire network may want to use BGP-based diversion. In contrast, the operator of a single machine (or a single service such as a Web site) needs to only divert traffic destined to one host and could thus use a DNS-based setup. A DNS-based setup is typically easier to configure and requires fewer resources (e.g., you do not need to configure and speak BGP). A downside of DNS-based diversion is that it only works for proxiabale services and applications (e.g., Web sites).

2.5 Moment of Mitigation

Network traffic diversion can be done in an *on-demand* or *always-on* manner. In the case of *always-on* protection, traffic is *always* routed towards the DPS infrastructure, even if a customer is not under attack. Thus, if DNS-based

diversion is used, an address lookup always results in an IP address that routes to the DPS infrastructure. In the BGP case the DPS will never withdraw the customer's IP subnet announcement.

With *on-demand* protection, a DNS change or BGP prefix announcement is made in response to an attack, and negated when mitigation has completed. For the prior, the DNS change depends on the method by which DNS-based diversion is used:

- **IP Address Only** – The owner of a domain changes the IP address records from an IP address that does not route to the DPS infrastructure to a DPS-assigned IP address. Multiple address records may need to be changed if the domain has more than one. All changes can later be reverted to stop diverting traffic.
- **Canonical Name** – Since the DPS controls the authoritative name server for the canonical domain name, the DPS can make changes in a manner similar to that outlined above.
- **Name Server** – The DPS controls the authoritative name server for the protected domain name and as such it can change the IP address record(s) accordingly.

On-demand protection can be manual or automated. As an example of the latter consider customer-premise mitigation equipment (i.e., an in-line appliance) that sends out an alert to the DPS in case an attack is too large to handle in-line. In such a hybrid approach, the DPS can initiate *on-demand* protection automatically.

2.5.1 BGP Blackholing

BGP blackholing is a network traffic filtering mechanism that can be used to bring about Denial-of-Service attack mitigation. This operational countermeasure leverages the Border Gateway Protocol and uses the *communities attribute* [38], an extension to BGP that enables BGP speakers to tag prefix announcements with additional information [29]. In the context of blackholing, a specific blackholing tag (or set of tags) allows for one network (i.e., autonomous system) to request another network (e.g., an upstream provider) to drop, i.e., null-route, all traffic destined to the tagged prefix [10]. By dropping attack traffic closer to the source, i.e., before it even reaches the target network, the risk of congestion on interconnecting links as well as within the target network is reduced. An example blackholing tag is `asn:666`, in which `asn` marks the requestee-AS by which blackholing is requested (i.e., the blackholing provider),

(a) The victim requests 172.16.0.1/32 to be blackholed

(b) The provider drops all traffic destined to 172.16.0.1/32

Figure 2.3: BGP blackholing by an upstream provider

and 666 is a specific value that the blackholing provider recognizes for the purpose of blackholing requests.

Figure 2.3 shows a scenario in which blackholing is used to mitigate an attack on the IP address 172.16.0.1. The red lines indicate attack traffic coming from multiple sources on the Internet, and converging on the interconnecting link. To deal with this attack, the victim AS requests the prefix 172.16.0.1/32 to blackholed by the requestee-AS, AS1 (see Figure 2.3a). To this end, it tags the announcement with AS1:666. The requestee-AS recognizes this community and filters all traffic destined to the tagged prefix. Once traffic is filtered, the risk of congestion within the requesting AS's network, as well as on the interconnecting link, is reduced (see Figure 2.3b).

Blackholing can also be implemented in Internet eXchange Points (IXPs) [30, 37], in which case blackholing can be requested through the IXP route server. Traffic can then be null-routed by IXP members at their points of ingress. Figure 2.4 shows a possible scenario. On the left-hand side the blackholing request is sent by the victim AS to the route server and next propagated to IXP members. Note that ASX in the tag is the AS number of the IXP. On the right-

(a) The victim AS requests (through the route server) other IXP members to blackhole 172.16.0.1/32

(b) The IXP members drop all traffic destined to 172.16.0.1/32 at their ingress switch interfaces

Figure 2.4: BGP blackholing within an Internet eXchange Point

hand side the IXP members have dropped, at their ingress switch interfaces, all traffic destined to 172.16.0.1/32.

As the previous examples have already shown, a tagged prefix can be as specific as a /32, i.e., specify a single IPv4 address. But it may be less specific and thus cover a range of IP addresses. Best practices suggest that prefixes less specific than /24 should not be blackholed [68]. It is also recommended to blackhole as specific as possible, in order to limit the impact of blackholing on adjacent IP space that is not under duress of DoS [63]. An AS can request for a blackhole to be removed either by re-announcing a prefix without a blackholing tag, or by withdrawing the prefix.

Blackholing is an attractive DoS mitigation technique for multiple reasons. First, it has the potential for quick activation, by just announcing a tagged prefix. Secondly, differently from cloud-based protection services, blackholing does not involve network traffic diversion towards third-party infrastructure. And thirdly, blackholing is relatively inexpensive in terms of operation. At the same time, blackholing can also be considered coarse-grained with respect to alternative mitigation solutions. This is because all network traffic to blackholed prefixes is dropped indiscriminately.

As a general technique to filter traffic, BGP blackholing is also at times used for reasons other than DDoS mitigation (e.g., censorship).

Attack Characterization

In this chapter we take the first steps towards studying the Denial-of-Service (DoS) phenomenon on an Internet-wide scale. We start by identifying viable data sources to report on the DoS problem. We then use said data to embark on a rigorous characterization of attacks and attack targets.

The data sources outlined in this chapter are also used in studies in later chapters of this thesis (i.e., Chapters 4, 5, 6, and 7). It is therefore recommended that this chapter is read first.

The study discussed in this chapter was previously published as part of the paper “Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem” in Proceedings of the 2017 ACM Internet Measurement Conference (IMC’17) [51].

3.1 Introduction

A rigorous characterization of the DoS phenomenon faces tremendous challenges. First of all, the need for sustained operational infrastructure to capture indicators of a variety of different types of DoS attacks poses a challenge. In addition, complex data fusion techniques are required to integrate heterogeneous raw data sources as well as meta-data to support classification and correlation of attack events.

In this chapter we take the first steps toward this goal. We start by identifying data sources that provide indicators of DoS attacks on an Internet-wide scale. We then systematically combine these data sources with supplemental

metadata to enable a macroscopic characterization of attacks and attack targets.

In this chapter we focus on answering the following questions:

- Which data sources on DoS activity can we use to broadly report on the scale and characteristics of the DoS problem?
- Can we realistically fuse, extract and correlate existing data sources on DoS activity to extract macroscopic as well as detailed insights about attacks?
- What does the DoS landscape look like in terms of attacks and attack targets?

3.2 Data Sources on DoS Activity

We identified two distinct data sources that provide global indicators of DoS activity. First, the UCSD Network Telescope, which captures evidence of DoS attacks that involve randomly and uniformly spoofed IP addresses. And second, AmpPot honeypots, which capture reflection and amplification DoS attacks – an attack type that involves specifically spoofed IP addresses. We detail these data sources in Section 3.2.1 and Section 3.2.2, respectively. It should be noted that these data sources recur in several later chapters, where we broaden our analyses to include, e.g., mitigation.

3.2.1 Randomly Spoofed Attacks

The UCSD Network Telescope [17] is a largely-unused yet routed /8 network operated by the University of California, San Diego. Network telescopes, which are also referred to as darknets, passively collect unsolicited traffic – resulting from scans, misconfigurations, bugs, and backscatter from Denial-of-Service attacks, etc. – sent to routed regions of the address space that do not contain any hosts.

Figure 3.1 shows how a network telescope can pick up backscatter from DoS attacks. The example attack shown is a SYN flood attack. In such an attack, traffic consists of TCP SYN packets, which involves the first packet type from a three-way TCP handshake. The source IP address in these packets is forged, i.e., set randomly to a spoofed IP address, by the attacking infrastructure. The victim may, provided that its link is not (immediately) saturated by the attack, upon receipt of a SYN packet, answer with a handshake response, i.e., TCP SYN|ACK. If the spoofed address is within the network telescope’s address space,

Figure 3.1: A bird’s-eye view on DoS activity inference using the UCSD Network Telescope.

the response packet will be sent to the telescope (rather than to the actual source of the attack packet), where the packet can be collected and analyzed.

We implemented the detection and classification methodology described by Moore et al. [80] to identify randomly spoofed Denial-of-Service attacks in the data collected at the telescope. The implementation comes as a Corsaro [61] plugin that we have also released publicly as open source [62]. Our plugin uses the same three-step processes described by Moore et al. First, we identify and extract backscatter packets. Second, we combine related packets into attack “flows”. Note that relatedness is based on the victim IP address. Finally, we perform attack classification and filtering.

The example in Figure 3.1 uses SYN|ACK backscatter packets, but we classify many more response packets as backscatter. Specifically, TCP SYN|ACK, TCP RST, ICMP Echo Reply, ICMP Destination Unreachable, ICMP Source Quench, ICMP Redirect, ICMP Time Exceeded, ICMP Parameter Problem, ICMP Timestamp Reply, ICMP Information Reply, or ICMP Address Mask Reply. We then aggregate such packets into flows based on the victim IP address (i.e., the source IP address of the backscatter packets), and we expire flows using the same conservative 300 second timeout described by Moore et al. In the final attack classification and filtering step, we compute statistics about the number of unique spoofed source IP addresses and the number of different ports used. We also compute four metrics of estimated attack intensity: (i) the overall number of packets; (ii) the overall number of bytes; (iii) the attack duration; and (iv) the maximum packet rate per second (pps_{max}). We use the

same conservative thresholds described by Moore et al. to filter low-intensity attacks, discarding those with: (i) fewer than 25 packets, (ii) a duration shorter than 60 seconds, and (iii) a maximum packet rate lower than 0.5 pps. A packet rate of 0.5 pps to the telescope corresponds to an estimated packet rate of 128 packets per second to the victim (the number should be multiplied by 256 to account for the telescope’s IP address space coverage). Assuming 1500 B packets, 0.5 pps observed then corresponds to an approximate attack traffic volume of 1.5 Mbps to the victim.

While the maximum packet rate can be used as an indicator of the attack intensity, this statistic also reflects the capability of the victim to endure the attack. That is, a high-intensity attack to a well-provisioned victim will likely result in a higher observed maximum packet rate than the same attack directed at a poorly-provisioned victim. The reason for this is that poorly-provisioned links are more likely to become saturated by a high-intensity attacks.

The UCSD Network Telescope covers approximately $1/256$ of the IPv4 address space. This means that any sizable attack, i.e., one that involves many packets with randomly and uniformly spoofed IP addresses, is likely to be visible on this darknet.

3.2.2 Reflection and Amplification Attacks

The second data source on attack activity is formed by reflection and amplification honeypots from the AmpPot project [65]. The novel and open-source AmpPot honeypot aims to track reflection and amplification DoS attacks by mimicking reflectors. To be appealing to attackers, AmpPot emulates several protocols known to be abused in reflection attacks. Specifically, the protocols QOTD, CharGen, DNS, NTP, SSDP, MSSQL, RIPv1, and TFTP. This way, AmpPot can be found by attackers scanning for reflectors and be “abused” in subsequent DoS attacks.

Figure 3.2 shows how AmpPot honeypots can log reflection and amplification attacks. The example shown uses a DNS reflection attack. In a reflection attack, the attacker sends requests allegedly coming from the victim. This is done by forging the source IP address of request packets to be the IP address of the victim. Upon receiving a forged request, the reflector typically sends its response (which is a DNS answer in the shown example) to the victim rather than to the actual source of the request. AmpPot pretends to be a usable reflector and logs forged requests as they arrive.

The AmpPot data also contains an attack intensity measure (rps_{avg}), expressed in terms of the average number of requests per second (e.g., DNS queries). As a honeypot is part of a larger group of amplifiers used during an attack, the attack intensity depends on the total number of amplifiers involved. While

it is unclear how many other amplifiers are involved in each attack, our best guess is that the total number of amplifiers will not vary significantly among attacks using the same amplification vector.

In order not to cause harm in actual attacks, AmpPot only replies to sources sending fewer than three packets per minute. However, logging the requests allows for various information about an attack to be inferred, including the IP address of the victim, the start and end of the attack, and also the request rate, which can be used as a measure of intensity. To distinguish attacks from other traffic (e.g., scans for reflectors), AmpPot only considers events exceeding 100 requests.

Figure 3.2: A bird's-eye view on DoS activity inference using a reflection and amplification honeypot.

Within the AmpPot project, an initial set of eight honeypots was installed in November 2014. The set has since been expanded to 24 honeypots. To prevent skew in the data set by either country or autonomous system, the honeypots are distributed both geographically¹, as well as logically, among various cloud providers and machines operated by volunteers. It has been shown that by making the honeypots attractive to attackers (in terms of the the amplification that attackers can achieve), 24 honeypot instances are sufficient to catch most reflection and amplification DoS attacks on the Internet [65].

¹ 11 honeypots are located in America, 8 in Europe, 4 in Asia and 1 in Australia.

3.2.3 Attack Coverage and Target Metadata

Many types of DoS attacks involve spoofed IP addresses, either for the purpose of hiding the attacking infrastructure, or to enable reflection attacks. As we have pointed out previously, any sizable DoS attack that involves randomly and uniformly spoofed IP addresses should be visible on the UCSD Network Telescope. Moreover, 24 honeypot instances catch most reflection and amplification attacks, which involve specifically spoofed IP addresses (i.e., that of the victim). The two previously outlined data sources on DoS activity therefore complement each other in terms of the attack types registered. It should be noted that attacks in which network traffic is sent to victims without spoofing (e.g., by botnets that do not bother to spoof the source IP addresses) are not covered by the two data sources.

Both data sources provide targeted IP addresses. These IP addresses can be augmented with metadata to study target characteristics (e.g., the target’s location). We use *NetAcuity Edge Premium Edition* data [39] to add geolocation information. And we use *Routeviews Prefix-to-AS mappings* data [13] to add BPG routing metadata.

3.3 Data Sets

In this chapter we analyze and correlate two data sets built from the previously identified data sources on DoS activity. Each data set covers a specific, recent two-year period (March 1, 2015 – February 28, 2017) and, again, contains attack events with different characteristics.

Table 3.1 summarizes both data sets. The network UCSD-NT data set has 12.47 M randomly spoofed attack events, involving 2.45 M unique targets (i.e., unique IP addresses). The AmpPot data set has 8.43 M reflection attacks, targeting 4.18 M unique targets. We will further discuss these data sets as we detail our study of attacks and targets, in Section 3.4 of this chapter.

source	#events	#targets	#/24s	#/16s	#ASNs
UCSD-NT	12.47 M	2.45 M	0.77 M	31057	25990
AmpPot	8.43 M	4.18 M	1.72 M	41678	24432
Combined	20.90 M	6.63 M	2.49 M	43041	32580

Table 3.1: DoS attack events data. We consider two years of data from the UCSD Network Telescope and from AmpPot honeypots to infer DoS attack events. Over the two years we observe more than 20 million events targeted at more than 2 million /24 network blocks.

3.4 Analysis of Attacks

In the following sections we will characterize attacks as well study various properties that relate to attack targets.

3.4.1 A third of the Internet attacked

Together, our data sets of attack events account for *20.90* M attacks, targeting *6.34* M unique IP addresses, over a two-year period (Table 3.1). We observe a total of *2.19* M unique /24 network blocks that host at least one target, which is about a third of the ~ 6.5 M /24 blocks recently estimated to be active on the Internet [95, 117]. For repeated attacks against the same IP address, we see fewer events per target IP in the AmpPot data than in the UCSD-NT data, which we attribute to more follow-up in randomly spoofed attacks. Combined numbers for both data sets also show overlap in targets, which we investigate further in this section.

3.4.2 Around 30 k DoS attacks a day are visible

Figure 3.3 shows statistics over time for the two years’ worth of attack events. The top graph shows randomly spoofed attacks, i.e., those in the UCSD-NT data set. The *attacks* curve shows the number of events seen on each day, which averages out to about *17.1* k daily. The *unique targets* curve is noticeably lower than the *attacks* curve, on each day, highlighting that some targets are hit more than once on the same day by randomly spoofed attacks.

The middle graph of Figure 3.3 shows statistics over time for attack events in the AmpPot data set. The average number of *attacks* is about *11.6* k daily. In this case, the *unique targets* and *attacks* curves are not as far apart as for randomly spoofed attacks, reflecting a lower average number of events per target IP address.

Finally, the bottom graph in the same figure shows the combination of attack events from both data sets. In total, we observe an average of *28.7* k attacks per day. The curve of *unique targets* is not the sum of the unique targets seen in each data set individually. This is because some targets are hit by both randomly spoofed and reflection DoS attacks on the same day, which we investigate in more depth in Section 3.4.10.

A takeaway from these results is that each day we see attacks on tens of thousands of unique target IP addresses, spread over thousands of autonomous systems, as shown by the *targeted ASNs* curves. The combined events as well as the individual time series also reveal spikes and plateaus in terms of the number

of attack events. We evaluate such outliers in Chapter 4, where we focus on the impact of attacks.

Figure 3.3: The number of attacks over time (black curves), and the number of IP addresses (gray curves), /16 network blocks (blue curves), and ASNs (orange curves) targeted over time for: randomly-spoofed DoS attacks observed in the UCSD-NT data set (top graph), attack events in the AmpPot data set (middle graph), and the union of these two data sets (bottom graph). Note that the combined data is not simply the sum of the top two graphs: in some cases we observe targets attacked by both randomly-spoofed, and reflected DoS attacks, on the same day.

country	#targets	%
US	625 k	25.56%
China	256 k	10.47%
Russia	140 k	5.72%
France	126 k	5.14%
Germany	103 k	4.20%
Other	1200 k	48.91%

(a) UCSD-NT

country	#targets	%
US	1232 k	29.50%
China	416 k	9.96%
France	323 k	7.73%
GB	266 k	6.37%
Germany	216 k	5.18%
Other	1727 k	41.26%

(b) AmpPot

Table 3.2: The targeted IP addresses and percentage of all observed attacks per-country (based on the NetAcuity Edge IP geolocation database). While this ranking mostly follows Internet space usage patterns, we find some notable exceptions, e.g., while Japan ranks *3rd* in recent address space usage estimates, it ranks *25th* and *14th* in the UCSD-NT and AmpPot data respectively. On the other hand, Russia and France rank higher in terms of attacks compared to address space usage.

3.4.3 By-country target ranking follows Internet space usage patterns, with some notable exceptions

We rank the most-commonly targeted countries, based on the geolocation metadata of target IP addresses. Table 3.2a shows that more than one fourth of randomly spoofed attack targets geolocate to the United States, with *25.56%* (or *625 k*) of all unique IP addresses. China follows second, with *10.47%* of targets. These two countries also rank first and second for reflection attacks in Table 3.2b, respectively with *29.5%* and *9.96%* of *4.18 M* unique target IP addresses. In general, we find that the two rankings are largely consistent and mostly reflect available statistics of Internet address space utilization (e.g., , routed space or estimated used space [35]). However, there are some notable exceptions. While in recent estimates Japan ranks third (*6.22%* and *6.33%* of space announced on BGP or inferred as actively used, respectively [8]), in the UCSD-NT and AmpPot data sets it ranks *25th* and *14th*, respectively. Russia and France, are instead examples of countries that in these attack data sets rank higher than in estimates of Internet space usage. In the case of France, we found out that this shift is mostly due to attacks to OVH, a large hoster that was heavily attacked in 2016 [90].

3.4.4 TCP is the preferred protocol in randomly spoofed attacks

The distribution of IP protocols in the attack events in the UCSD-NT data set provides an overview of the flooding approach used. Table 3.3 shows that the majority of these attacks involve TCP (79.4%), while UDP and ICMP follow at 15.9% and 4.5%, respectively. ICMP in this distribution denotes ICMP attack traffic (e.g., a ping flood, which leads to ICMP Echo Reply backscatter). In case an ICMP Destination Unreachable message reaches the telescope, we register the protocol of the quoted packet, e.g., UDP for a UDP packet that could not reach its destination. Other protocols, which include, for example, IGMP, account for 0.2% of attack events.

IP protocol	TCP	UDP	ICMP	Other
events (%)	79.4%	15.9%	4.5%	0.2%

Table 3.3: IP protocol distribution. The percentage of all attacks per IP protocol as observed in the UCSD-NT data.

3.4.5 NTP is the preferred reflector protocol in reflection and amplification attacks

The AmpPot data set does not suggest which specific service was targeted by reflection attacks. Instead, we observe which amplification vector (i.e., reflector protocol) was used by the attacker. Table 3.4 shows a distribution of the protocols chosen by attackers. NTP leads with 3.38 M attack events, accounting for 40.08% of the 8.43 M reflection attacks seen over two years (Table 3.1). The second and third placed, DNS and CharGen, account for 26.17% and 22.37%, respectively. Examples of protocols following SSDP and RIPv1 in terms of occurrence are MS SQL and TFTP.

NTP is also the most-used protocol for reflection according to various vendor reports. While we find similarities between our results and vendor reports, we also find differences. As vendor reports are based on customer-specific data and oftentimes do not state the scientific method used, we do not delve into these similarities and differences further.

3.4.6 Randomly spoofed attacks tend to last longer. 10% last more than an hour and a half

Each target is attacked a certain amount of time. Attacks typically last minutes up to hours. Figure 3.4 shows the distributions of the attack duration in our

type	#events	%
NTP	3.38 M	40.08%
DNS	2.21 M	26.17%
CharGen	1.89 M	22.37%
SSDP	0.71 M	8.38%
RIPv1	0.23 M	2.27%
Other	0.01 M	0.73%

Table 3.4: Reflection protocol distribution. Number of attacks (and percentage of all attacks) per reflection protocol as observed in the AmpPot data.

Figure 3.4: Duration of attacks. The distributions of duration in the UCSD-NT (top graph) and AmpPot (bottom graph) data sets.

data sets. The top and bottom graphs refer to randomly spoofed and reflection attacks, respectively. About 40% of randomly spoofed attacks last five minutes or shorter. Attacks in the UCSD-NT data set last at least one minute due to the minimum duration threshold that we outlined in Section 3.2.1. We find that roughly the top 10% of randomly spoofed attacks last 1.5 hours or longer. While attacks in the UCSD-NT data set can last longer than a day, these cases are rather scarce ($\sim 0.2\%$). The mean duration is 48 minutes and the median is 454 seconds.

For attack events in the AmpPot data set we find that 50% of attacks last 255 seconds or shorter. The top 10% of attacks last 40 minutes or longer, and roughly 6% of attacks last an hour or longer. The mean attack duration is 18 minutes and the median duration is 255 seconds. We note that because of how the honeypots operate, they cap attack event durations at 24 hours. As only $\sim 0.02\%$ of attacks last 24 hours we don't expect this cap to significantly affect the results.

3.4.7 More than a thousand attacks of medium to maximum intensity occur on a daily basis

The attack data sets contain intensity attributes, which we use to analyze the strength of attacks. For randomly spoofed attacks we see the maximum number of packets per second reaching the network telescope during the attack (pps_{max}). This rate can range from tens to tens of millions of packets per second. To infer an estimate of the packet rate reaching the victim, assuming the attack is using uniformly random spoofing, the rate should still be multiplied by 256 (see Section 3.2.1). For reflection attacks we observe the average number of requests made to the reflector per second (rps_{avg}). This number can range from below one to hundreds of thousands. The reason for the comparative difference in the higher ranges is because reflection attacks are amplified, and need fewer packets to reach large traffic volumes.

We use these attributes to estimate the intensity distributions over attacks. Figure 3.5 shows the results for attack events in the UCSD-NT data set. A steep curve shows that about 70% of attacks generate only about 2 pps_{max} reaching the telescope, which translates to an estimated attack rate of 512 packets per second to the victim. For about 17% of the attacks, the telescope observes more than 10 packets per second (an estimated attack rate of 2560 packets per second to the victim). The mean and median values are 107 and 1, respectively.

For attacks in the AmpPot data set, given that the total number of amplifiers will not vary significantly among attacks using the same amplification vector (see Section 3.2.2), we analyze the intensity distribution separately per protocol. Figure 3.6 shows the overall distribution for all attack events, as well as separate curves for the top five used reflector protocols. Note that these five reflector protocols are involved in all but 10 k attack events, as shown in Table 3.4. For most protocols, about 70-90% of attacks see a gradual increase in the number of requests per second (rps_{avg}), starting as low as below one on average, to a couple thousand. The number of requests involved clearly varies per protocol. Taking NTP as an example, roughly the first 90% of attacks see up to 2000 packets per second, whereas the top intensities involve tens to hundreds of thousands of packets per second. These distributions are also different

Figure 3.5: The intensity distribution for attack events in the UCSD-NT data set. The number of packets per second (max) should be multiplied by 256 to estimate the the packet rate reaching the victim.

compared to the UCSD-NT data, which we attribute to the different nature of attack events. The overall mean and median values are 413 and 77 requests per second, respectively.

Figure 3.6: The intensity distribution for attack events in the AmpPot data set. We show the distribution for the top five reflector protocols used, as well as the overall distribution.

Figure 3.7 shows attack events that have a medium intensity or higher, over time, for both data sets combined. We consider an attack event to be of medium intensity or higher if its intensity is at least the mean of all intensities in the corresponding data set. On average, daily, we observe 1.4 k attacks within this intensity range, compared to the overall average of 28.7 k attacks per day (Figure 3.3). We study one of the peaks visible in the curve in Section 4.3.

Figure 3.7: High-intensity attack events over time. The number of attacks with a medium or higher intensity, per day, for the UCSD-NT and AmpPot data sets combined.

type	#events	%
single-port	7.56 M	60.6%
multi-port	4.91 M	39.4%

Table 3.5: Number of target ports distribution. Number of attacks (and percentage of all attacks) per target port cardinality in the UCSD-NT data.

3.4.8 Web, gaming, and MySQL ports are the most attacked in randomly spoofed attacks

Randomly spoofed traffic sent to flood a victim can target one or multiple ports. One reason to target a single port is because the attacker wants to take down a specific networked daemon. Another reason is because the port is known (or assumed) not to be filtered by a firewall. Table 3.5 shows, for the 12.47 M attack events of this type, the number of events that targeted strictly one port (60.6%), as well as those that involved multiple ports (39.4%). Note that for the AmpPot data we do not make a port number distinction, because we do not keep track of the typically ephemeral target port in the reflected packet.

We map the ports of attacks that target only a single port to applications, i.e., services, on the basis of both IANA port assignments, as well as commonly used port numbers. Table 3.6 shows the results of this mapping for TCP and UDP. We show in Figure 3.6 per protocol the top five potentially targeted services, along with their share of the distribution within that respective protocol. We say “potentially” because we do not know if the service was listening at the

type	#events	%	type	#events	%
HTTP	2.83 M	48.68%	27015	225.4 k	18.54%
HTTPS	1.20 M	20.68%	37547	24.8 k	2.04%
MySQL	0.06 M	1.12%	32124	17.1 k	1.41%
DNS	0.06 M	1.07%	28183	16.9 k	1.39%
VPN PPTP	0.06 M	0.99%	MySQL	15.8 k	1.30%
Other	1.60 M	27.46%	Other	916 k	75.32%

(a) TCP

(b) UDP

Table 3.6: The distribution of target ports in the UCSD-NT data set. We show the top five potentially targeted services – based on IANA port assignments – for randomly spoofed attacks to a single port using TCP (left) and UDP (right).

time of the attack. Moreover, the port might have been chosen by an attacker merely to penetrate a firewall to perform a service agnostic attack.

Table 3.6a shows the results for TCP. HTTP ranks first with 2.83 M attack events, which account for 48.68% of 5.81 M single target port attacks on TCP. HTTPS ranks second with 20.68%. The third place goes to MySQL (3306/TCP), with a share of 1.12%, which is significantly lower than HTTP(S). For UDP, in Table 3.6b, the most-attacked port is associated with various on-line multiplayer games and the Steam platform.² About 75% of the UDP attack events target ports that do not rank among the top five, which is because these attacks are spread out over the roughly 65 k remaining ports.³

There are two important takeaways from these results. First, while attacks associated with on-line gaming are most apparent for UDP, most other attack events for UDP are spread out over the full port range. Second, more than two thirds of all attack events over TCP potentially target Web infrastructure (69.36%).

3.4.9 Randomly spoofed attacks against Web ports are more intense

Given the prominent presence of Web ports (i.e., 80 & 443) in the UCSD-NT data set we evaluate the mean and median intensity of attacks that potentially target Web ports. We find that the mean (maximum per attack) rate observed at the telescope is 226 packets per second – corresponding to an estimate of almost 60 k packets per second. This is a change upward from 107 for all

²<http://steampowered.com/>

³A few examples of services over UDP that follow the fifth placed MySQL are NTP (123/UDP) and NetBIOS (138/UDP).

randomly spoofed attacks (the median remains the same). We also compared the duration statistics with their overall counterparts and find that the mean drops to 10 minutes (down from 48) and the median drops to 240 seconds (down from 454). We thus find that attack events that involve Web ports are more intense than the overall, while lasting shorter.

The prevalence of strong attacks on Web infrastructure ports prompts us to study the impact of attacks on Web sites in more detail. We will do this in Chapter 4, where we focus on the impact of attacks. Attacks on Web sites may also trigger the outsourcing of protection to a DDoS Protection Service, which we study in Chapter 5.

3.4.10 Randomly spoofed and reflection and amplification attacks are sometimes used jointly against the same target

Finally, we study cases in which targeted IP addresses show up in both the UCSD-NT and the AmpPot data sets. That is, the targets are hit by randomly spoofed attacks as well as reflection attacks over time. The UCSD-NT and AmpPot data sets have 282 k unique target IP addresses in common (Table 3.1). Out of 282 k targets, 137 k were hit simultaneously by joint attacks, i.e., attacks that overlap in time. An example of a joint attack is a SYN flood combined with an NTP reflection attack. The vast majority (77.1%) of randomly spoofed attacks co-participating in attacking a victim involve a single port in the UCSD-NT data set: we see an increase from 60.6% (Table 3.5), suggesting that joint attacks are more likely to target a specific service. The target port distribution of randomly spoofed attacks jointly involved with reflection attacks has more attacks to 27015/UDP (53% up from 18.54%), which suggests that joint attacks might be used to gain an edge in on-line gaming. For TCP, an increase in HTTP from 48.68% to 50.23% is seen. While the latter is a subtle change, it could indicate that serious attackers, i.e., those who launch both randomly spoofed and reflection attacks, target Web services more often.

The distribution of IP protocols in randomly spoofed joint attacks is similar to that of all randomly spoofed attacks and shifts only by tens of percents. For reflection attacks co-participating in attacking a victim, we find that CharGen's use drops by half, to 11.5%, while the other four protocols in the top five gain. NTP gains most with an increase to 47.0%.

The autonomous system most-commonly targeted by joint attacks is *AS12276 (OVH)*, with 12.3% of 137 k unique joint attack targets. *China Telecom* is placed second with 5.4%. *China Unicom's AS4837* is third (3.1%). When considering joint attacks, the per-country distribution does not differ significantly from those for single attacks.

The first-most and second-most countries to which joint targets geolocate are the US and China, with 24.4% and 20.4% , respectively. France comes third (9.5%) and Germany fourth (6.5%). These four countries are also in both top fives in Table 3.2a and Table 3.2b, and in the same order. Russia, which was not in the top five for reflection attacks, is fifth placed for joint attacks (4.1%).

3.5 Related Work

We consider as related work efforts to characterize DoS attacks in general. Such characterizations include, for example, target properties (e.g., geolocation), traffic characteristics (e.g., protocols used), and attacker properties (e.g., malware fingerprinting).

In 2006, Moore et al. [80] characterized DoS attacks by analyzing events inferred from backscatter packets to a large network telescope. The authors analyze 22 traces of 1-2 weeks each, captured between 2001-2004, totalling 68.7 k events. We incorporated their methodology in our work. Their initial trace is 14 years older than our UCSD-NT data set. Comparing results, ours show that the DoS landscape has since changed. As an example, while still dominant, TCP's presence in randomly spoofed attacks has reduced. Moreover, we find a prevalence of single-port attacks.

Krämer et al. [65] and Thomas et al. [103] both present a characterization of attacks from events captured in a set of amplification honeypots. While in both papers the focus is more on reflection attacks in general, in this paper we focus on the correlation with randomly spoofed attacks and on target characteristics. A different view on DoS attacks is given by Santanna et al. [99], who study traffic and source characteristics of the attacks generated on-demand by means of a set of 14 booters. Differently from our paper, this research focuses on the attackers (i.e., the misused infrastructure).

To our knowledge, the last study to characterize DoS attacks at scale by combining multiple, independent data sets dates back to 2006, when Moa et al. used three data sets [73] in their work. Two data sets came from anomaly detection systems and a third was inferred from backscatter. Their analysis covers 35 k attack events, measured over a month, which does not compare in scale with our study. The authors find a *TCP* preference similar to Moore et al., using the same methodology.

More recently, in 2015, Wang et al. [112] analyzed a set of 51 k attack events derived from botnet Command & Control (C&C). Their data set covers a seven-month period and accounts for attacks launched using 674 botnets of 23 different botnet families. They too find joint attacks, in their case by different botnet

instances. Furthermore, they show that Web services (i.e., *HTTP*) are the preferred target of many attacks.

The industry regularly releases reports that characterize attacks and trends [21, 28, 40, 76]. However, these reports are based on customer-specific data, and oftentimes do not state the scientific method used.

3.6 Concluding Remarks

This chapter presents our first steps towards rigorously characterizing the DoS phenomenon on an global scale. We identified two data sources that provide Internet-wide indicators of DoS activity. Specifically, a large network telescope, and honeypots instrumented to observe reflection and amplification attacks. We systematically fused and correlated these diverse data sources on DoS activity, and augmented the attacks data with metadata such as BGP prefix-to-AS mappings and IP geolocation.

Our results speak to the scale of the DoS problem. About a third of all /24 networks recently estimated to be active on the Internet were involved in attacks over a two-year period. We observed roughly 30 k attacks each day, a number higher than previously reported. Our successful fusing of diverse data sources revealed characteristics of multiple attack types launched jointly that are not apparent from any single data source.

While most of the measurement infrastructure used for the work in this chapter was pre-existing, a significant challenge was posed by data fusion, extraction, correlation and visualization. The experience in overcoming this challenge resulted in our capability to extract macroscopic as well as detailed insights about DoS attacks, the results of which we presented in this chapter. Equally importantly, it paved the way for studies presented in later chapters of this thesis.

Impact of Attacks

The previous chapter presents a macroscopic characterization of attacks as well as attack targets. A sensible curiosity that follows is: what is the potential impact of attacks? Our analysis revealed that more than two thirds of TCP-based attacks target Web infrastructure ports. We also found a prevalence of stronger attacks on these ports. This prompted us to study the potential impact of attacks using Web sites as a measure. The results are presented in this chapter.

The study discussed in this chapter was previously published as part of the paper “Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem” in Proceedings of the 2017 ACM Internet Measurement Conference (IMC’17) [51].

4.1 Introduction

In the previous chapter we joined diverse data sources that provide indicators of DoS activity. We revealed well over 20 million attacks over a two-year period. This comes down to around 30 k DoS attacks on average each day, targeting tens of thousands of unique target IP addresses, spread over thousands of autonomous systems. Naturally, this prompts the question of how much damage attacks can do, if successful.

In this chapter we build upon the previous chapter by studying the potential impact of attacks. We choose Web sites as a measure as in the previous chapter we discovered that Web servers are a prominent attack target.

In this chapter we focus on answering the following questions:

- Which data source can we use to study the potential impact of attacks on prominent attack targets (i.e., Web sites)?
- What is the potential impact of attacks on the Web, if successful?
- Are Web infrastructure targets more likely to be hit by Web specific protocols and ports?
- How do Web hosters factor into the potential impact of attacks?

4.2 DNS Measurement Data

The UCSD-NT and AmpPot data sets – previously identified and used in Chapter 3 – contain per attack event the IP address of the attacked target. To evaluate the potential impact of attacks using Web sites as a measure we need a historical mapping between IP addresses and Web sites hosted. To obtain this mapping we use active DNS measurement data from the OpenINTEL project [20, 107].

OpenINTEL is a large-scale, active DNS measurement platform that collects daily snapshots of the content of the DNS. It builds snapshots by structurally querying all the domain names under a full zone, i.e., Top-Level Domain (TLD), a set of Resource Records (RRs). OpenINTEL covers a large number of TLDs and the resulting data notably includes domain name to IP address mappings (i.e., A records).

Figure 4.1: A bird-eye’s view of the OpenINTEL measurement and analysis architecture.

We are among the founders of OpenINTEL and have been actively involved in its development and operation since the beginning. Our work on OpenINTEL

has paved the way for contributions made by this thesis, as well as by other research efforts. Unsurprisingly, OpenINTEL provides a data source that recurs in many chapters of this thesis. We also often rely on the OpenINTEL Hadoop infrastructure to perform analyses.

We provide a glimpse at the full OpenINTEL architecture here. Figure 4.1 shows the various stages of OpenINTEL. Stage I relates to zone (i.e., TLD) collection. Stage II relates to the daily measurement. And stage III relates to data storage and analysis.

In this chapter we identify Web sites that are potentially affected by attacks by looking for **A** records on **www** labels that, at the time of a given attack, mapped, i.e., resolved, to the attacked IP address. We take the presence of a **www** label in the DNS as an indicator that Web content was present (or intended) at the time of the attack. We did not probe each and every domain name to see if Web content was present.

start	#days	source	#Web sites	#data points	size
2015-03	731	.com	173.7 M	1045.9 G	23.5 TiB
		.net	21.6 M	121.0 G	2.8 TiB
		.org	14.7 M	90.7 G	2.1 TiB
		Combined	210.0 M	1257.6 G	28.4 TiB

Table 4.1: Active DNS data set. We use two years of DNS data collected by the OpenINTEL platform to infer Web sites and associated IP addresses for the .com, .net, and .org gTLDs. In this data set we find 210 M domains that we classify as Web sites (i.e., those with a **www** label).

We use a subset of the TLDs that OpenINTEL measures. Specifically, we use DNS data for the three generic TLDs (gTLDs) .com, .net, and .org, which combinedly cover roughly 50% of the global domain namespace [15]. Table 4.1 shows the details of the data set. For each of the three gTLDs, we show the total number of Web sites over the two-year period. For example, for .com (the largest TLD), a total of 173.7 million Web sites were seen. The *data points* column shows the total number of collected data points, examples of which are **CNAME** and **A** RRs. The total number of data points is 1.258 trillion. The *size* column shows the size of the compressed measurement data using Apache Parquet columnar storage [18], with a total of 28.4 TiB. On the last day of the studied, two-year period the three gTLDs account for 153 million active **www** domain names.

We combine the OpenINTEL data set with DoS attack events from UCSD-NT and AmpPot data. In fact, we use the same attacks data sets as in our previous chapter on attack characterization. This provides an opportunity to

source	#events	#targets	#/24s	#/16s	#ASNs
UCSD-NT	12.47 M	2.45 M	0.77 M	31057	25990
AmpPot	8.43 M	4.18 M	1.72 M	41678	24432
Combined	20.90 M	6.34 M	2.19 M	43041	32580

Table 4.2: DoS attack events data. We consider two years of UCSD-NT and AmpPot data.

compare properties of attacks that specifically involve targets with Web site associations with overall attack properties, for which we will occasionally be making backreferences to Chapter 3. All three data sets span the same two-year period (March 1, 2015 – February 28, 2017). Table 4.2 summarizes the attacks data sets again for convenience.

4.3 The Effect of Attacks on the Web

In this section we evaluate the potential effect of attack events on the Web. We consider the subset of attack events that target IP addresses for which we can determine Web site associations, using the active DNS measurement data set described previously, in Section 4.2. We find Web site associations on 572 k of the 6.34 M unique target IP addresses in the attack events. This means that of uniquely targeted IP addresses, at least 9% host one or more Web sites.

While analyzing Web site associations we may find that multiple Web sites share an attacked IP address. As a consequence, an attack on a single IP can potentially affect millions of Web sites simultaneously. These cases occur when an IP address is used by a larger party, such as a hoster. In case of multiple associations, a single Web site as well as the hoster as a whole may have been the intended target of the attack. Regardless, all Web sites that share that IP address can potentially be affected. We identify large parties by looking at routing information for the attacked IP address, by looking at a common name server in the NS record, or a common CNAME through which Web sites expand to the shared IP address. To elaborate the last point: in some cases a CNAME record in the DNS can reveal more about a Web site than the Web site’s IP address. For example, some hosters rely on Amazon AWS, which means that IP routing information points to Amazon and not to the hoster. A customer-specific CNAME that all Web sites share might still reveal the hoster.

Figure 4.2: Web site associations with IP addresses targeted by attacks. Each bar indicates the number of unique target IP addresses (y axis) associated with a number of Web sites within a given bin (x axis).

4.3.1 Many target IP addresses belong to large hosters, with each mapping up to millions of Web sites

Figure 4.2 shows the number of Web sites affected by attack events. Each bar, i.e., bin, represents a “co-hosting” group, which indicates how many Web sites were associated with a targeted IP address at the time of an attack. The magnitude of each group is the number of target IP addresses within the group.¹ More than a third of these IP addresses (~ 211 k) were associated with a single Web site at the time of an attack, whereas, at the other end of the distribution, 169 targets hosted 1 M to 3.6 M Web sites potentially affected by the attack event (3.6 M is the maximum in the right-most group in the graph). We note that this maximum is found on a target IP that is routed by DOSarrest, one of many DDoS Protection Services that we consider in other chapters of this thesis. Google and Amazon are other examples with (various) IP addresses in this group.

The active DNS data set used in this chapter accounts for Web sites under .com, .net, and .org. Our estimate of Web sites per target IP address is therefore a lower bound. IP addresses could be associated with Web sites in TLDs that we have not considered (e.g., .tk). For this reason the inclusion of additional TLD data may introduce previously unmatched IP addresses into the distribution. Moreover, IP addresses that we have currently mapped to one “co-hosting” group could be shifted to a larger group as the result of additional

¹Each IP address can contribute once to a “co-hosting” group in this visualization.

mappings. As a result, the overall distribution might change. To analyze this possibility we considered `com`, `net`, and `org` individually. The shape of Figure 4.2 is similar for the three individual distributions, which suggests that the distribution among “co-hosting” groups would not drastically change even if we were to consider additional TLD data.

4.3.2 Isolating Web targets reveals an even more pronounced majority of TCP-based randomly spoofed attacks and NTP-based reflection attacks

Recall from Section 3.2.1 that the UCSD-NT data provides us with attack transport layer protocol and port information. Considering the UCSD-NT data set, we find that randomly spoofed attacks against IP addresses that are associated with Web sites primarily use TCP and target Web infrastructure ports. Specifically, 93.4% of attacks use TCP and 87.60% of attacks target either port 80/TCP or 443/TCP. In the previous chapter, in which we studied attacks independently of Web site associations, we found that 79.4% of all randomly spoofed attacks use TCP (see Table 3.3). We also found that 69.36% of all attacks target Web infrastructure ports (see Table 3.6a). Upward shifts are apparent, which more strongly suggest that Web infrastructure indeed is being targeted.

Considering reflection attacks in the AmpPot data set, we find that NTP is the most commonly used reflector type on targets with Web site associations, in 54.69% of cases, up from 40.08% (see Table 3.4).

4.3.3 Over two years, 64% of inferred Web sites were hosted on IP addresses targeted by attacks

Figure 4.3 shows, for every day in our two-year observation period, the total number of Web sites associated with attacked target IP addresses on that day. The top graph is for all attack events, and the bottom one is for attack events with a medium to high intensity. In each graph, the gray curve shows the number of Web sites (potentially) affected, in millions, whereas the black curve shows the (smoothed) percentage that the involved Web sites make up of all inferred Web sites in the measured namespace, meaning `com`, `net` and `org` (right y axis).²

We link almost 134 M unique Web sites to all attack events observed over the two-year period. This comes down to about two thirds of all inferred Web sites (see Table 4.1). The average number of attack-associated Web sites is just under 4 M per day, which corresponds to about 3% of all inferred Web sites in

²Note that for smoothing we interpolate a cubic spline between the median number of affected Web sites per month.

Figure 4.3: Web site associations with attacked targets over time. The number of Web sites on attacked IP addresses for all attacks (top graph) and medium to high intensity attacks (bottom graph). The left y axis shows the number of Web sites and the right y axis shows the percentage of all Web sites in the measured namespace.

the measured namespace. Note that multi-day attacks, i.e., those that cross day boundaries, count only towards the day on which the attack was started. The daily average number of Web sites associated with attacks of medium intensity or higher is 1.7 M (1.3%).³ The fraction of Web sites that are potentially affected daily is considerable, which does not come as a surprise given the large number of ASNs and /24 prefixes that we see attacked daily (see Section 3.4).

4.3.4 Attacks on large hosting providers

In the number of affected Web sites, various peaks are discernible, the largest of which involves 11.82% of all Web sites. We evaluate this peak, along with three others, as examples of the potential effect of attacks on the Web.

The first peak (see Φ in Figure 4.3) on March 12, 2015 involves attacks that associate with a little over 15 million Web sites, which is 11.82% of all inferred Web sites. We identified several large hosters as attack targets on this

³We consider an attack to be of medium intensity if its intensity is the mean of all intensities in its attack events data set.

day. A significant number of Web sites associated with attacks were hosted by *GoDaddy*, through a set of about twenty targeted IP addresses, all routed to their AS. Moreover, a large number of Web sites was associated with *WordPress*, primarily through two consecutive IP addresses that belong to *Automattic Inc.*, the company behind *WordPress*. Another IP address routed to the security infrastructure of CenturyLink, one of the DPS providers considered in this work, which shows that the attack was probably mitigated. Many of the target addresses appear as joint attacks in the honeypot and the telescope data sets, with low to medium intensities.

The second peak (see ②) on October 10, 2015, involves 11.7 million Web sites. Among the targets we find several large hosters such as *Squarespace* and *OVH*. Another prominent target is a domain names reseller that is hosted in *Amazon AWS*.⁴ The third peak we investigate, occurs on November 4, 2016 (see ③). It involves a little over 13 M Web sites. About 10 M of these Web sites are hit by an attack of high intensity, as can be seen in the bottom graph. This number is largely made up by *GoDaddy*-hosted Web sites. A significant number is also associated with *Wix.com*, a Web site development platform.⁵ *Squarespace* is among the targets. The final example (see ④) is for February 25, 2017. This peak involves 14.1 million Web sites, hosted by various companies, such as *GoDaddy*, *OVH*, *Network Solutions*, and a variety of hosting companies that are subsidiaries of the *Endurance International Group (EIG)*.

Overall, the three most frequently attacked larger parties that we identify over the two-year period are, in order, *GoDaddy*, *Google Cloud*, and *Wix*. Other names include *Squarespace*, *Gandi*, and *OVH*.

We encountered during our analysis several IP addresses that can be linked to the mail infrastructure of a large number of domain names. In these cases it is not a domain name's *www* label that maps to an attacked IP address, but rather its mail exchanger record (*MX*). For example, we found that *GoDaddy's* e-mail servers, which are used by tens of millions of domain names, are frequently targeted by DoS attacks. These findings suggest that DoS attacks may also have a significant impact on mail infrastructure. As we use Web sites as a measure to study the potential impact of attacks we do not explore this further at this point, but do stress it as possible future research direction.

⁴This company has its own AWS *CNAME*, which allowed us to identify it even though the IP belongs to AWS.

⁵Wix hosts in AWS, but uses Incapsula for DDoS mitigation, which is something we previously outlined in [57].

4.4 Related Work

We consider as related work efforts to measure the effects of DoS attacks. Welzel et al. measured the impact of botnet attacks by monitoring for targets in botnet C&C [115]. Their study covers 646 unique targets, acquired from 14 botnet instances of two botnet families (*DirtJumper* and *Yoddos*). Following attack commands, the authors systematically measure the victims for adverse effects. Occasionally they find that the IP address of a Web site changes following an attack, e.g., in an attempt to mitigate, by pointing it to *localhost*. In a few cases the IP address change is made to (quote) “professional load balancing and DDoS protection services,” but this is not investigated further.

Noroozian et al. [84] study the consequences of victimization patterns in targets of DDoS-as-a-Service (e.g., booters). Their focus is on the demographics of the target population. Their results show, among others, that most of the victims are users in access networks, and that the number of attacks in broadband ISP is proportional to the number of ISP subscribers. Similarly to us, their study is also based on two years of AmpPot data. However, we focus on capturing a larger spectrum of attack events by correlating amplification honeypots data with network telescope data.

In terms of effects at a higher level, a DoS attack can have financial consequences for businesses, which could face an increase in security costs, or a loss of customers following an attack [113]. While DDoS intensity peaked at 400 Gbps [93] in 2014 and to 600 Gbps in early 2016 [60], the race to the largest DDoS has already reached 1 Tbps in late 2016 with the attack against the hosting company OVH [90]. However, it is not only about how heavy the hammer is, it is also about what it might break. The DDoS attack performed by the Mirai botnet against the service and DNS provider Dyn [44] has provoked a cascading effect that prevented East Coast users from accessing services such as Twitter, Spotify, or Reddit.

4.5 Concluding Remarks

This chapter presents an additional step towards characterizing the DoS phenomenon. We built on our previous chapter on attack characterization and shifted our focus to the potential impact of attacks. Driven by previously having found Web servers to be the most prominent attack target, we chose to use Web sites as a measure, creating a need for a mapping between attacked IP addresses and Web sites. We identified another diverse data source to this end, the OpenINTEL project, and successfully fused and correlated DNS data from this source with our data on DoS attacks.

Our results speak to the potential widespread impact of attacks. On average – for **com**, **net** and **org** – we found 3% of Web sites involved daily (1.3% for stronger attacks). Nearly two-third of all inferred Web sites were associated with at least one attack over a two-year period. In absolute numbers, this adds up to well over a 130 million Web sites. It is worth noting that the considered namespace ‘only’ accounts for about 50% of the global DNS namespace. We also showed that bringing down a single IP address can potentially take millions of Web sites offline simultaneously. All our results stress the importance of mitigation, which we will investigate further in later chapters of this thesis.

DDoS Protection Services

At this point we shift our focus from attacks to mitigation solutions. We start with DDoS Protection Services, which amount to the first of two global mitigation solutions that we study in this thesis (the other being BGP blackholing). In this chapter we study the adoption of protection services on the Internet, and we look at factors that influence adoption.

In this chapter we devise a new data source – on the use of protection services. We fuse it with data sources identified in previous chapters. The work in Chapter 7 involves data on the use of protection services as well, which is why we recommend that this chapter is read before Chapter 7.

*This chapter is based on two previously published papers. First, “Measuring the Adoption of DDoS Protection Services” in *Proceedings of the 2016 ACM Internet Measurement Conference (IMC’16)* [57]. And second, on a specific part of [51].*

5.1 Introduction

Distributed Denial-of-Service attacks have steadily gained in popularity, their intensity ranging from mere nuisance to severe. Our preceding chapters underpinned the scale of the DDoS problem: *(i)* in terms of attack occurrence we found 30 k attacks daily; and *(ii)* in terms of potential impact we found more than 130 million Web sites to have been involved in attacks (i.e., hosted on attacked infrastructure) over a period of a few years – millions of which are hit by higher intensity attacks every day.

The loss of revenue for targets of attacks has given rise to a market for mitigation solutions, among which DDoS Protection Service (DPS) providers, to whom (potential) targets can outsource protection.

Protection services, which can be contracted to protect Web sites, among others, are the first global mitigation strategy that we study in this thesis. In this chapter we investigate the adoption of cloud-based protection services worldwide. We focus on nine leading providers and make our outlook on adoption on the basis of active DNS measurements. We introduce a methodology that allows us, for a given domain name, to determine if traffic diversion to a DPS is in effect. Our methodology also allows us to distinguish various methods of traffic diversion and protection.

In this chapter we focus on answering the following questions:

- Which data do we need to study the use of DDoS Protection Services?
- What does adoption on the Internet look like over time, for leading commercial providers?
- In what manner do end-users use protection services?
- How dynamic is the use of protection services?
- Which factors drive DPS adoption?

5.2 Data Sources on DPS Use

To study the adoption of protection services we need to first identify a data source that enables us to infer the use of such services. Then, for the purpose of having a well-defined scope, we need to make a decision on which providers to study. The following two sections deal with these topics.

5.2.1 Network Traffic Diversion

In our background chapter, Chapter 2, we explained that the use of protection services involves using the DNS or BGP to divert network traffic. The OpenINTEL project measures the DNS records on which various DNS-based diversion mechanisms rely. This allows us to devise a methodology to infer DNS-based diversion from OpenINTEL data. In particular, we can infer DPS use from **A**, **CNAME** and **NS** records (see Section 2.3.2).

To infer BGP-based network traffic diversion we evidently need to consider BGP routing information. To this end we supplement IP address records in OpenINTEL data with autonomous system numbers. We do this analogously

to how we augmented attack target IP addresses in UCSD-NT and AmpPot attacks data (see Section 3.2.3).

5.2.2 Leading Providers

We will focus on leading providers of cloud-based protection services. We make our selection of providers based on a 2015 *Forrester Wave* report [47]. Forrester is an advisory market research company that follows a publicly available methodology. We select all nine providers from the report. Specifically, Akamai, CenturyLink, CloudFlare, DOSarrest, F5 Networks, Incapsula, Level 3, Neustar, and Verisign.

5.3 Methodology and Data set

Now that we have explained which data sources we will use and which protection services we selected for our outlook on adoption, we will outline our methodology and describe the resulting data set.

We will use the Hadoop ecosystem to analyze OpenINTEL data to infer if and how domains are protected by DDoS Protection Services. The various steps of our methodology, as well as the resulting data set, are described next.

5.3.1 Inferring the use of DDoS Protection Services

We analyze OpenINTEL data to infer the use of DPS providers by domain names. Several of the selected providers offer DNS-based traffic diversion and (optionally) authoritative name server protection. More specifically, we detect CNAME-based redirection by checking whether the CNAME expansion of x contains a DPS reference. Similarly, the NS record of x will reference a DPS if the DNS zone of x is managed by that DPS. Lastly, the ASN of x 's IP address(es) can also reference a DPS.¹ By doing this longitudinally, our analysis reveals, per day, if domain x uses one (or several) of these methods.

We detect DPS references in CNAME and NS records based on the second-level domain (SLD) contained therein. For example, we found that Incapsula uses the SLD *incapdns.net* in CNAME records. To identify SLD and ASN references, we take the following steps:

1. Given a DPS name, we infer a candidate set of AS numbers for the DPS by searching AS-to-name data by name

¹We use *Routeviews Prefix-to-AS mappings* [13] to add BGP routing metadata to the IP addresses that OpenINTEL observes in the DNS.

2. By analyzing domain name to IP address mappings in OpenINTEL data, as well as the associated AS numbers (from BGP routing metadata), we find a set of domain names that reference the given DPS by ASN (as inferred in step 1)
3. For the domain names that result we next find frequently occurring SLDs in **CNAME** and **NS** records (e.g., *incapdns.net*)
4. We use the resulting SLDs as **CNAME** and **NS** references for the DPS
5. Now using the **CNAME** and **NS** references as starting points, we check if we missed any AS numbers in the first step, or if we need to drop AS numbers that are not used by the DPS for mitigation purposes

Based on combinations of references and non-references we can analyze not only if, but also how domain names use a DPS. Take for example a domain that references a DPS by **CNAME** as well as by ASN, but not by **NS** record. This combination of references shows us three things. First, it shows us that the domain uses a **CNAME** record to give the DPS control over IP address records. Secondly, the ASN reference shows us that network traffic is actively being diverted. And thirdly, we learn that the DNS zone of this domain has not been delegated to the DPS.

By evaluating combinations of references we also identify frequently-used third parties, such as third-party name servers that are authoritative for large numbers of domains that switch on or off protection simultaneously.

5.3.2 Always-on and On-demand Use

To analyze if domain names use a DPS in an *always-on* or *on-demand* manner (see Section 2.5), we track DPS use per domain name on a day by day basis. If a given domain name references a DPS by ASN without gaps, we infer *always-on* use. We infer *on-demand* use if there are gaps. In the case of *on-demand* use, **CNAME**, **NS**, and **ASN** (non-)references reveal specifically how traffic diversion was effected. For example, if a domain name switches back and forth between two IP addresses, of which only one references a DPS by ASN, we infer DNS-based traffic diversion. We infer BGP-based diversion if the IP address does not change, whilst its ASN supplement does change between a DPS reference and a non-reference.

5.3.3 Data Set

We use 1.5 years worth of DNS measurement data for the generic TLDs (gTLDs) **.com**, **.net**, and **.org**. In addition, we use six months of data for the country-

code TLD (ccTLD) `.nl`, as well as for the Alexa Top 1 million. Table 5.1 details the OpenINTEL-provided data set. The column `#SLDs` shows the number of unique SLDs observed over the studied period. `#DPs` is the number of collected data points (i.e., `CNAME`, `A`, `AAAA`, and `NS` measurements). The `size` column shows the compressed measurement data size in the OpenINTEL Hadoop cluster using Parquet columnar storage [18] (before replication). The three gTLDs contain about 50% of the global domain namespace; on the last day of the data set they contain a little over 152 million names.

Source	start	days	#SLDs	#DPs	size
<code>.com</code>	2015-03	550	161.2M	534.5G	17.5TiB
<code>.net</code>	2015-03	550	20.2M	62.4G	2.1TiB
<code>.org</code>	2015-03	550	13.8M	46.7G	1.5TiB
<code>.nl</code>	2016-03	184	5.9M	10.4G	2.1TiB
Alexa 1M	2016-03	184	2.2M	1.7G	77.5GiB
Total			203.3M	655.7G	23.3TiB

Table 5.1: Active DNS measurement data set

Table 5.2 shows the ASN and SLD references for the selected DPS providers, obtained using the steps described in Section 5.3.1. We note that for some of the studied providers the references can overlap with customers of other services. For example, for Akamai domain names that do not use *Kona Site Defender* (their reverse proxy) and *Prolexic Routed/Connect* (their BGP-based mitigation solution) can be traced to the found AS references. Some providers do not work with `CNAME` redirection, but through delegation can *change* the IP address of a domain (e.g., Verisign’s *Managed DNS service*). Some providers (e.g., F5 Networks & DOSarrest) do not offer DNS-based options.

5.4 Adoption and Characteristics of Use

5.4.1 General Overview

Using the references in Table 5.2, we analyze the three main gTLDs and find per day the number of domains that use the DPS providers under consideration. We consider use by domains on their second level, meaning that multiple references in the DNS zone of a domain are counted as one. Figure 5.1 shows how the number of distinct SLDs varies over time. The figure is dominated by many “anomalous” peaks and troughs, which can involve millions of domains. For example, the peak on the 5th of March, 2015 involves about 1.1 M domain names. The anomalous trend that is apparent in the largest gTLD, `.com`, is

Provider	AS number(s)	CNAME SLD(s)	NS SLD(s)
Akamai	20940, 16625, 32787	<i>akamaiedge.net,</i> <i>edgekey.net,</i> <i>edgesuite.net,</i> <i>akamai.net</i>	<i>akam.net,</i> <i>akamai.net,</i> <i>akamaiedge.net</i>
CenturyLink	209, 3561	—	<i>savvis.net,</i> <i>savvisdirect.net,</i> <i>qwest.net,</i> <i>centurytel.net,</i> <i>centurylink.net</i>
CloudFlare	13335	<i>cloudflare.net</i>	<i>cloudflare.com</i>
DOSarrest	19324	—	—
F5 Networks	55002	—	—
Incapsula	19551	<i>incapdns.net</i>	<i>incapsecuredns.net</i>
Level 3	3549, 3356, 11213, 10753	—	<i>l3.net, level3.net</i>
Neustar	7786, 12008, 19905	<i>ultradns.net</i>	<i>ultradns.*</i>
Verisign	26415, 30060	—	<i>verisigndns.com</i>

Table 5.2: DDoS Protection Service references

replicated in *.net* and *.org*, which indicates that the anomalous behavior is transversal to the zones. Many of the larger anomalies are part of *on-demand* behavior, which we discuss in more detail in Section 5.4.4.

Figure 5.1: DPS use and zone breakdown

Figure 5.2 shows over time per DPS the number of domains that use any of the DPS’s services (the top line). As can be seen, some of the larger anomalies can be traced to Incapsula (e.g., the previously mentioned peak in March 2015

in Figure 5.1). Some providers show very few anomalies, and contain more domains than the more anomalous providers on their “quiet” days. For example, CloudFlare versus Incapsula in March 2015, were it not for the anomalous peak.

Figure 5.2: DPS use per provider and protection method breakdown

Figure 5.3: DPS use and gTLD distribution over namespace

In Figure 5.3 we show the (average) distribution of the three main gTLDs over the roughly 50% of the global domain namespace that they cover, as well as the distribution of DPS using domains among these TLDs. Both distributions are remarkably similar, suggesting that there is no correlation between a zone and subscribing to a DPS.

5.4.2 Overall Growth

For our growth analysis we do not count anomalous peaks and troughs. We smooth shorter and smaller anomalies out by taking the median reference count over a time window of several weeks, while the large anomalies are cleaned manually. This way we largely separate *always-on* from *on-demand* use. Figure 5.4 shows the combined growth of the nine providers relative to the start of our data set, in about 50% of the global domain namespace. The overall expansion of the zones involved is also shown. A trend in the adoption of DPSs becomes apparent, which is largely driven by CloudFlare, DOSarrest, Incapsula, and Verisign (see Figure 5.2). Other providers such as F5 Networks and CenturyLink contribute to incidental decrease (e.g., the dip in March 2016). As shown, DPS use has grown by $1.24\times$ over 1.5 years, which exceeds the overall expansion of $1.09\times$, from about 140 M to 152 M domains.

Figure 5.4: Growth of DPS use in 50% of the DNS

We applied the same procedure to our six-month data for `.nl` and the Alexa Top 1 million. Figure 5.5 shows the results. A growth trend of 10.5% against 1.8% is apparent for `.nl`. For Alexa (which has a fixed size) the growth in DPS use is 11.8% .

Figure 5.5: Growth of DPS use in `.nl` and Alexa

5.4.3 Protection Methods

As outlined in our background chapter (see Section 2.3.2), various ways exist to use the DNS to divert network traffic to a DPS. Our reference analysis has shown that a DPS may support more than one way (see Table 5.2). Figure 5.2 shows per DPS a breakdown of the various use cases (the marked NS, CNAME, and AS curves).

Difference in use among the nine DPS providers can be discerned, even where providers support the same DNS use cases. For example, if we look at the use of delegation, we find that CloudFlare’s authoritative name servers are used significantly, by about 75% of CloudFlare-using domains on average (compare CloudFlare’s NS and overall, top line). For Incapsula, however, only about 0.02% of domains use delegation, i.e., are using the *Incapsula NS Protection* service (this NS line is not visible). Verisign sits somewhere in the middle. During most of the first eleven months (March 2015 until February 2016), the number of Verisign-using domains that used delegation (i.e., their *Managed DNS* service) was even higher than those that diverted traffic (compare Verisign’s NS and AS curves). We suspect that the predominant use of CloudFlare’s *Authoritative DNS* among CloudFlare customers is because the service is free. We analyzed for a single day the set of full names of CloudFlare’s authoritative name servers, most of which are given a male or female name, followed by `.ns.` and *cloud-*

flare.com, the NS SLD reference. There are 403 such names on April 30th, 2016, with *kate.ns.cloudflare.com* the most-referenced (by 112 k domains).

5.4.4 Dynamic Behavior

Third-party Anomalies

We have traced many of the larger anomalies shown for each DPS to *on-demand* or *always-on* use by third parties, and in one case to a DNS issue at a third party. A few examples will follow. For **Incapsula**, Web site development platform *Wix* causes repeated swings of millions of domain names², such as the peak in April 2016 (see ④) that involves 1.76 M names. A second anomalous example for Incapsula is the increase in June 2016 (see ②), which we traced to “an opportunistic private equity fund around Internet domain names.”³ Specifically, this increase of about 170 k domain names can be traced to SiteMatrix (a domainer). Most of **Verisign**’s larger anomalies can be traced to *ENOM* (a registrar) and *ZOHO* (a hosting party), accounting for changes of up to 700 k domains.⁴ The February 2016 anomaly for **CloudFlare** (see ③) involves ~247 k *Namecheap*-hosted domains.⁵ The anomalous trough on November 22nd, 2015 for **Akamai** (see ⑥) was caused by ~716 k domains that can be traced back to Sedo Domain Parking. We infer that this was a DNS issue at Sedo, since the number of measured domains with a *sedoparking.com* NS SLD also dipped that same day. Our final example is the significant drop of domains in February 2016 for **CenturyLink** (see ⑤). We traced this to a platform that offers “Expert tools to manage domain registration, sales and monetization.”⁶ Some of the observed anomalies involve multiple providers. For example, the March 2015 peak for Incapsula has an opposing trough in F5 Networks (see ⑥ & ⑦).⁷

Daily Fluctuations and Repeated Anomalies

To study if repeated anomalies involve the same set of domain names, we analyzed the daily flux per provider in terms of first seen and last seen domain

²Wix domains normally route to Amazon AWS (AS14618) through a *amazonaws.com* CNAME. During diversion, Wix name servers answer A records in various Wix-owned prefixes that are announced by Incapsula.

³<http://www.sitematrix.com>

⁴Several ENOM-owned /24s route to Verisign (AS26415) during diversion, and to ENOM (AS21740) normally. Similar for ZOHO, with two prefixes normally in AS2639.

⁵The domains share a Namecheap NS SLD (i.e., *registrar-servers.com*) that answers CloudFlare-announced addresses.

⁶Here, a Fabulous-owned name server, starts giving A answers for ~355 k domains that previously routed to two prefixes announced by CenturyLink’s AS3561.

⁷Here, two Wix-owned prefixes switch back and forth from F5 Network’s AS55002 to Incapsula’s AS19551.

Figure 5.6: Flux of DPS use per provider

names. This way, if protection is turned on and off several times for a set of names, the names involved will contribute to influx at most once, and to out-flux at most once. Figure 5.6 shows per DPS the delta of first seen and last seen counts, grouped in two-week time windows. As shown, repeated anomalies in Figure 5.2 can be traced to the same sets of domain names.⁸ For example, the large influx for Incapsula in March 2015 indicates that many of the same domains were involved in the anomalous plateau that starts in May 2015. A second takeaway is that over time some providers contribute more gradually to DPS adoption than others, of which CloudFlare is a prime example, since its influx is rather spread out.

On-demand Use

Our outline of some of the larger anomalies shows that many can be traced to *on-demand* use, while some we suspect are *always-on* domains because of only an

⁸Time grouping and variations in the customer base of third parties can change the flux magnitudes somewhat.

Figure 5.7: On-demand peak duration occurrences

upward or downward edge. Since our measurement period is finite we cannot easily determine if an opposing edge can be found outside the measurement period. Moreover, a domain that shows a single period of use, i.e., peak, could either be a short-lived *always-on* customer, or brief *on-demand* use. Thus, it is not trivial to classify the type of DPS use. To gain more insight into dynamic behavior among the various providers we estimate for each a set of *on-demand* domains, which is done on the basis that the domains show at least three peaks over 1.5 years. For the sets of domain names, we analyzed the peak durations in days over the 1.5 year period. Figure 5.7 shows the results as the CDF of peak occurrences. For providers that show signs of highly anomalous behavior from day to day, the majority of peak occurrences are short-lived (i.e., $P(\text{duration} \leq \text{days}) = 0.8$). A good example is Neustar, with 80% of all peaks lasting four days or fewer, which we suspect is because their *always-on* solution is a hybrid in which traffic is not continuously diverted to the cloud.⁹

⁹<https://www.neustar.biz/resources/faqs/ddos-faqs>

5.5 Attack Effects on Adoption

Now that we have a better understanding of how protection services are used and who drives adoption, a logical question to ask next is to which extent having been under attack influences DPS adoption. Web site owners who maintain their own hosting, as well as hosting companies that provide hosting infrastructure on a larger scale, may start outsourcing protection to a DPS after being targeted by a DoS attack. Intuitively, a causal link between attacks and DPS adoption exists. To the best of our knowledge, however, there is no work that addresses this at scale. In this section we study whether attacks on Web sites have an effect on protection service adoption, and to which extent.

From our data on DPS use we can analyze if, and when, Web sites adopted a DPS. In the previous chapter we created a link between attacks and Web sites. If we fuse these data sets we can see if adoption follows one or more attacks. We refer to this process as *migration*.

5.5.1 Integrating Attacks Data with DPS Use

provider	#Web sites
Akamai	5.86 M
CenturyLink	0.87 M
CloudFlare	4.27 M
DOSarrest	7.04 M
F5	3.58 M
Incapsula	3.78 M
Level 3	0.47 M
Neustar	10.78 M
Verisign	4.34 M
VirtualRoad	< 100

Table 5.3: DDoS Protection Service use. For each of the 10 DPS providers that we consider, we identify the Web sites they provide protection services for.

At this point we thus revisit data sets already used in preceding chapters: (i) the UCSD-NT and AmpPot attacks data sets (see Section 3.3); and (ii) the mappings of attacked IP addresses to 210 M Web sites (see Section 4.2). In terms of DPS use, we will focus on Web sites in the three larger gTLDs (i.e., .com, .net and .org). Whereas earlier in this chapter we studied 1.5 years of DPS use, we will consider two years from here on out to match the observation period of the other data sets (March 1, 2015 – February 28, 2017). As before, we consider the leading providers Akamai, CenturyLink, CloudFlare, DOSarrest,

F5 Networks, Incapsula, Level3, Neustar, and Verisign. We add a tenth provider, VirtualRoad, which is a non-commercial provider that protects Web sites run by journalists, activists, and human rights workers. While VirtualRoad is a small provider that does not have a major impact on adoption, by including it we consider in our analysis also attack targets that would not normally outsource protection to a commercial DPS.

Table 5.3 shows the details of the (extended) DPS use data set in terms of the total number of Web sites that we associate with each of the ten providers, over two years. This data set tells us, for all the Web sites inferred, the day of migration to the DPS in question, provided that day is within the observation period.

5.5.2 Taxonomy of Web sites in DPS

Figure 5.8: Web site taxonomy. Nodes are annotated with the estimated number of web sites in each category (and the percentage of the parent category population). The root of the tree represents the overall set of domains (over the two years we study) that we infer to be Web sites (i.e., those with a `www` label). We find that of these 210 M Web sites, 64% were hosted on attacked IP addresses (at the time of an attack) at least once during our two year observation period.

We define a classification taxonomy for Web sites according to the tree in Figure 5.8. The root of the tree represents the overall set of domains (over our two year observation period) that we infer to be Web sites (210 M). We then split this set into two: those for which we observed attacks (134 M), and those for which we did not (76 M). We find that the majority of Web

sites (64%) were observed to be on attacked IP addresses over the course of two years. Then, for both of these categories (**attack observed** and **no attack observed**), we identify those Web sites that either already use a DPS (**preexisting customers**) – either from the beginning of our data set, or the first time they are found in the DNS – and those that do not (**non-preexisting customers**). We find a much higher percentage of **preexisting customers** in domains for which we observed attacks (24.9 M, 18.6%) than for those where we did not observe attacks (0.67 M, 0.89%), suggesting that Web sites we observe to be attacked during our two year observation period may have been previously attacked. Finally, the bottom level of the tree identifies those Web sites that were **non-preexisting customers**, but either did migrate (**migrating**) or did not migrate (**non-migrating**) to using a DPS. In the case of **attack observed** Web sites, we consider a Web site to be **migrating** if it is found in the DPS data set *after* we observed it being under attack. For **no attack observed** Web sites, we consider it to be **migrating** if it is found in the DPS data set *after* it is first seen in the DNS. While we do find a slightly higher percentage of Web sites migrate after an attack (4.7 M, 4.31%) compared to those that migrate even when no attack is observed (2.5 M, 3.32%), it should be noted that since we do not observe all attacks, the **no attack observed** migrations may still have been influenced by an attack. We also find the percentage of Web sites that either already used a DPS, or during our study migrated to using a DPS, to be much larger for those Web sites that were attacked (22.1%) compared to those for which we did not observe an attack (4.2%).

While our list of 10 protection services is not exhaustive – AWS (Amazon) and GHS (Google) actually offer DoS protection that we cannot infer and, therefore, the many Web sites they host count towards **non-migrating** in our classification – we take this into account in the following analysis. Additionally, because our attack events and DPS data sets cover the same time range, it is possible that we incorrectly classify attacks that occur close to the beginning and/or end of our observation period. More specifically, attacks that overlap the beginning of our DPS data set may have already prompted migration, thus resulting in an incorrect **preexisting customer** classification; similarly, attacks starting near the end may result in migration after our observation period, thus causing in an incorrect **non-migrating** classification. By shortening the observation period of the attacks data by one month on either end and repeating our analyses, we verified that these potential misclassifications have a negligible effect on the overall Web site class distribution.

We manually checked a small sample of Web sites to gain insight into the types of Web sites that are among various combinations of hosting size groups

and customer classes.¹⁰ We sampled from the smallest (i.e., $n = 1$), as well as the largest (i.e., $n \geq 10^6$) hosting groups, and for each of the three DPS customer classes (i.e., the leaves of the **attacked** subtree in Figure 5.8). For the largest hosting group, among those that migrate to a DPS after an attack, we find many Web sites that can be traced to the Wix Web site development platform. The Web sites we visited have either personal or business content. Among **non-migrating** within the largest hosting group, we find a lot of landing pages that can be traced to a domain reseller that uses AWS for hosting, as well as personal and business Web sites hosted in Google Cloud. Among the **preexisting customers** we find both personal pages and commercial Web sites such as a Web shop. For the smallest hosting group, we find among **migrating** and **preexisting customers** Web sites that belong to businesses, community Web sites (e.g., related to gaming), and occasionally content for a foundation. In one case we visited a Web site with radical right content, which may or may not speak to why the Web site was attacked. For the **non-migrating** class we find, among others, adult Web sites for (video) chat.

5.5.3 Repeated attacks are not a determining factor for migration

We observe a significant fraction ($\sim 14\%$) of Web sites attacked more than once within our observation period. We investigated if the number of attacks experienced by a Web site correlates with migration to a DPS. The top graph in Figure 5.9 shows the CDF for the distribution of all attacked Web sites as a function of the attack frequency: 7.65% of these sites are attacked more than 5 times. The bottom graph in the same figure shows instead the CDF, as a function of the attack frequency, for Web sites that migrate to a DPS after an attack event. In this case, the fraction of Web sites that were attacked more than 5 times is 2.17%. The comparison between the two distributions, suggests that being subject to multiple attacks is not a significant factor in subsequent migration to a DPS.

5.5.4 Earlier migration follows attacks of higher intensity

DoS attacks that severely affect Web sites are likely to create an urgency to mitigate. This notion makes it reasonable to assume that Web site owners (or hosters) who opt to outsource protection to a DPS will want to do so in an urgent manner. Table 5.4 shows the normalized attack intensity distribution over attacked Web sites. In the case a Web site is associated with multiple or

¹⁰We did not automatically verify for each potentially affected Web site if content was being served at the time of an attack.

Figure 5.9: The distributions of attack frequency for all Web sites (top graph) and those that migrate to a DPS following an observed attack event (bottom graph), a comparison of which suggests that being subject to multiple attacks is not a significant factor in subsequent migration to a DPS.

even simultaneous attacks (e.g., a target IP that appears both in the UCSD-NT and AmpPot data sets), we pick the highest normalized intensity value.

Intensity ($\leq$)	0.0	0.07	0.13	0.52	0.85	1.0
Web sites (%)	11.1	95.0	97.5	99.0	99.9	100.0

Table 5.4: Attack intensity distribution over Web sites. For select percentiles we show the normalized attack intensity in the UCSD-NT and AmpPot data sets. In case of joint attacks, we take the highest intensity.

Figure 5.10 shows the cumulative distribution functions of days it took Web sites to migrate, respectively for Web sites attacked with *any* intensity (slowest CDF), and with intensities in the 95-th, 99-th, 99.9-th percentiles of the normalized attack intensity distribution (Table 5.4). Comparing these CDFs highlights a drastic reduction of the latency between an attack and the effected site migrating to a DPS: almost all (98.6%) the top 0.1% Web sites by attack intensity transition to a DPS within 6 days, whereas for the top 1%, 5% and *overall* Web sites only 77.1%, 67.1% and 29.9% of them respectively transition within the same number of days. When considering the Web sites that transition to a DPS within a day from the attack, the difference between the top 0.1% class and the overall distribution is even more striking: 80.7% versus 23.2%, respectively. Differently from the number of attacks, the intensity of a

Figure 5.10: Migration delay for attack intensities. For various percentiles of the normalized attack intensity distribution, ranging from *any* to the *99-th*, we show the number of days it took for Web sites to migrate to a DPS. An urgency to migrate becomes apparent with increasing attack intensity.

DoS event strongly correlates with migration to a DPS, specifically in terms of speed, which intuitively suggests a sense of urgency in mitigating DoS damage and risks.

Large hosters can potentially skew the mitigation delay distribution by migrating many Web sites at once: if multiple Web sites are associated with an given attack of a given intensity, each Web site counts towards the CDF. We investigated this potential for skew and found that few migrating Web sites in the top *97.5-th* percentile were hosted in large numbers.

5.5.5 Attack duration does not strongly correlate with migration

Here we evaluate if attack duration may influence transition to a DPS and specifically timing. A target that is brought down by a successful attack will slow down or altogether stop backscattering packets to the telescope (see also Section 3.2.1). As such, attacks successful enough to trigger migration might be registered with shorter than actual durations in the UCSD-NT data set. Amplifiers on the other hand will still receive packets to reflect to the target, and thus have a better sense of the actual attack duration. For these reasons we only consider the durations from the AmpPot data set in this analysis.

Overall, we find that the number of days it takes migrating Web sites to migrate does not necessarily keep decreasing with an increasing attack duration,

Figure 5.11: DPS migration delay for longer attacks. The number of days it took for Web sites to migrate to a DPS following attacks with a duration ≥ 4 hours.

unlike is the case for attack intensity. Attacks longer than four hours in duration, which is the top 1% durations of all AmpPot events (see Table 3.3), lead to the smallest migration delay for **migrating Web sites**. Figure 5.11 shows the CDF for Web sites affected by attacks within this duration class: of all Web sites associated with attacks that last over four hours, 67.64% take a day or less to migrate, and 76% migrate within at most five days. However, more than half of the Web sites that migrate on the next day, following a 4 hours or longer attack, have a common denominator. Specifically, 482 k out of 800 k Web sites trace back to *Wix.com*, who starts outsourcing protection to Incapsula during our observation period. About 18% take two weeks or longer to migrate, suggesting that duration by itself is not always the deciding factor. We also find common denominators for longer migration delays. For example, 130 k Web sites hosted by *eNom* take more than three months (101 days) to appear as **migrating Web sites** (of Verisign).

Finally, we find larger parties that skew the results in favor of, as well as against, short migration delays. Comparing the two previous examples, the first one involves an attack that is three times as intense. Specifically, it involves a normalized attack intensity of 0.18 in the UCSD-NT data. This target appeared simultaneously in both the UCSD-NT and the AmpPot data sets. This finding leads us to conclude that in this case intensity rather than duration was the deciding factor.

5.6 Concluding Remarks

In this chapter we first studied the adoption of protection services and revealed that adoption has grown significantly on the Internet. Our results highlight a relative growth in adoption (among domain names) of $1.24\times$ over a recent period of 1.5 years. This growth surpasses an overall expansion of $1.09\times$ of the considered namespace – i.e., `.com`, `.net`, and `.org` – that represents 50% of the global domain namespace. Our results also highlight adoption trends in the ccTLD `.nl`, as well as among Web sites on the Alexa Top 1 million list. Respectively, we found relative growths of $1.11\times$ and $1.12\times$, over a period of six months.

Our methodology to infer DPS use can be used to analyze how domains divert traffic to a DPS, and whether or not optional services (e.g., name server protection) are in use. In our results we reveal differences in use of protection methods among the considered providers, even in cases where the compared providers support similar services. For some providers, only a small percentage of domains use delegation, which potentially leaves a part of a domain's DNS infrastructure (i.e., the authoritative name server) susceptible to DDoS attacks. We will study potential issues with DPS use in a later chapter in this thesis.

Our results also show that a large contribution to the user base and adoption of DPS providers is made by third parties, examples of which are Web hosts and domainers. Some of these larger players activate or deactivate DDoS protection for millions of domains from one day to the next, either by leveraging the DNS to divert traffic, or by having the DPS announce one or multiple IP prefixes.

In the second part of this chapter we studied to which extent having been under attack influences DPS adoption. To this end we fused our data sources on attacks and DPS use and correlated attack duration, repetition and intensity with adoption. While attack repetition and duration did not significantly influence migration, we unveiled an increased urgency for targets to migrate following attacks high in intensity.

BGP Blackholing

In this chapter we study the operational aspects of BGP blackholing at scale. Blackholing is the second of two global mitigation solutions that we study in this thesis. We put an emphasis on how blackholing is used when attacks occur, as well as what kind of attacks are mitigated with this measure. Additionally, as blackholing can be considered a self-inflicted Denial-of-Service of sorts, we will study the extent to which common Internet services become “collateral damage” as a result of blackholing.

This chapter is largely based on the previously published paper “A First Joint Look at DoS Attacks and BGP Blackholing in the Wild” in Proceedings of the 2018 ACM Internet Measurement Conference (IMC’18) [52].

6.1 Introduction

Results from earlier chapters in this thesis underpin the scale of the DDoS problem. We repeated this at the start of the previous chapter, in which we investigate DDoS Protection Services. And we repeat it again here, as we shift our focus to BGP blackholing, the second global mitigation strategy that we study in this thesis.

As explained in our background chapter, BGP blackholing is an operational countermeasure that builds upon the capabilities of the Border Gateway Protocol (BGP) to achieve DoS mitigation (see Section 2.5.1) [10]. It allows network operators to request an upstream provider (i.e., an ISP) to filter network traffic

destined to a network or even a single host. Blackholing is an appealing mitigation solution because it is relatively inexpensive in terms of deployment and operation. However, at the same time it is coarse-grained in the sense that all traffic to a blackholed prefix is dropped indiscriminately. Operators can thus choose to sacrifice the reachability of hosts to protect the overall network and its interconnecting links. In case blackholing cuts off a host that runs a service, legitimate users are denied access to that service as long as the blackhole remains in place. This brings about a Denial-of-Service in its own respect.

Although empirical evidence of blackholing activity on the Internet and in Internet eXchange Points (IXPs) is documented in literature [37, 42], a clear understanding of how BGP blackholing is used in practice when attacks occur is still missing. As blackholing is a coarse-grained approach to mitigation, one could imagine that it is used only to mitigate large attacks and only as a last resort if other, fine-grained solutions (e.g., protection services) are no longer applicable. As we will show in this chapter, this is not the case, which raises the question of what is the minimal effort needed by an attacker to trigger such a drastic countermeasure.

While arguably operators bring about a self-inflicted DoS by choice when they use blackholing, the inherent drawbacks of this mitigation technique raise the question: what is the collateral damage of such actions? Related work on this particular topic is often of an experimental or simulation-based nature [89, 110], and an Internet-wide quantification is, again, missing.

The goal of this chapter is twofold. First, we provide a first joint look at DoS attacks and BGP blackholing at an Internet-wide scale. And second, we quantify the drawback of blackholing. We focus on answering the following questions:

- What are the operational aspects of BGP blackholing following attacks in terms of, for example, mitigation response times?
- What can we learn about attacks that are mitigated through a measure as extreme as blackholing?
- Can we quantify the drawback of blackholing in terms of common Internet services that become cutoff from (legitimate users on) the Internet?

To the best of our knowledge, the work in this chapter amounts to the first large-scale empirical observation of DoS events and corresponding blackholing mitigation, as well as the first characterization of service collateral damage in the wild.

6.2 Methodology and Data Sources

We will analyze and study diverse data in this chapter to answer the aforementioned questions. This involves by now familiar data sources on, for example, attack activity, as well as data new and specific to this chapter. Most data sets for this chapter span a period of three years, which allows us to present a longitudinal overview of operational deployment and effects of blackholing. The following sections describe the data used as well as the methodologies behind newly created data.

6.2.1 DoS Attacks

We rely on the UCSD-NT and AmpPot data sources (see Chapter 3) to create a three-year data set of attack events, starting March 2015. Table 6.1 summarizes the data set in terms of attack events, targets and the number of target autonomous systems involved. We find *28.14* M attacks in total, targeting *8.58* M unique IP addresses. Recall that we observe a combination of multiple attack types as the two data sources are complementary. This allows us to identify *447.6* k instances of joint attacks, involving *176* k unique target IP addresses.

source	#events	#targets	#ASNs
UCSD-NT	15.89 M	2.94 M	29750
AmpPot	12.25 M	6.03 M	28425
Combined	28.14 M	8.58 M	36939
Joint	447.6 k	0.18 M	9218

Table 6.1: Denial-of-Service data from UCSD-NT and AmpPot. We find *28.14* M attacks, targeting *8.58* M unique IP addresses.

6.2.2 Blackholing Events

We obtain a data set of inferred blackholing events from publicly available BGP routing data, using a custom, extensible measurement system, implemented on the basis of the methodology described by Giotsas et al. [42]. Our blackholing data set covers a three-year observation period, starting March 2015 like the attacks data. The following sections will explain how we create the data set.

Public BGP data

We infer a data set of blackholing events from BGP routing data, using data from two projects that offer data publicly: (i) University of Oregon’s *RouteViews*

Project (RV) [16]; and *(ii)* RIPE NCC’s *Routing Information Service (RIS)* [11]. Both these projects collect, store and offer for download Internet routing data collected from globally dispersed collectors that each peer with one or multiple routers. While Packet Clearing House (PCH) is also a major source of public BGP routing data, we do not use this source. There are two reasons for this: *(i)* the *BGPStream* framework has limited support for PCH data (specifically, the data broker does not index PCH data); and *(ii)* the lack of support creates an interoperability problem with our reactive measurements (for reasons to become clear in Section 6.2.4).

Blackholing Communities

Within the BGP data, we look for BGP announcements tagged with a community that is likely to signal a blackholing request. Giotsas et al. [42] created a dictionary of such communities by applying natural language processing to resources where blackholing communities are likely to be documented. For example, they scanned Internet Routing Registry (IRR) records as well as network provider Web sites for terms such as “blackhole” and “null route”. And they validated many communities manually, for example through communication with operators. We use a copy of this dictionary, which provides us with 288 `asn:value` community tags, for 251 blackholing providers, using 74 distinct values (e.g., 666). We point out that the dictionary contains many BGP blackholing communities that are used in practice, but it is not necessarily complete due to methodological limitations [42].

Inferring Blackholing Events

We built a measurement system to infer blackholing activity, using the methodology of Giotsas et al. as a starting point. Our system uses *pyBGPStream*, an interface to the *BGPStream* framework for BGP data analysis [87]. We infer blackholing events by analyzing BGP updates and consider prefixes with a specificity of /24 to /32. Less specific prefixes are not commonly blackholed [42, 68]. We infer blackholing activity incrementally, by analyzing BGP updates. We do not parse routing tables for the beginning of the observation period. Consequently, we will miss blackholing events that were triggered before March 2015, which would anyway not be within the bounds of the attacks data set against which we match. Our measurement system is extensible and offers hook to, e.g., trigger reactive measurements. We defer discussing this functionality until Section 6.2.4.

To infer a data set of blackholing events we analyze public BGP data from 36 collectors.¹ Each blackholing event in the data set contains, most notably: (i) the blackholed prefix; (ii) a start time (i.e., activation time); (iii) an (optional) end time (i.e., deactivation time); (iv) the matched blackholing communities; and (v) the set of collectors on which prefix-related activity was observed.

More than one collector may see BGP activity related to the same prefix. We consider prefix-related activity to be part of one blackholing event if BGP record timestamps (partially) overlap in time. In such cases, we use timestamp extrema to determine the activation and deactivation time of the blackholing event. A blackholing event can be activated through a prefix announcement with a blackholing community set, and deactivated either through re-announcement without a blackholing community set (i.e., implicit deactivation), or through a prefix withdrawal (i.e., explicit deactivation). We presume consistent propagation characteristics between BGP announcements and withdrawals.

collectors	#events	#prefixes	#origins	#AS paths
34	1.30 M	146193	2682	31493

Table 6.2: Blackholing data set inferred from public BGP data. We infer 1.3 M blackholing events, involving 146193 prefixes.

Table 6.2 summarizes our data set. We find a total of 1.3 million blackholing events, involving about 146 k unique prefixes, and about 2.7 k ASNs from which the blackholing requests originated. 34 of the 36 collectors we consider see at least one blackholing event in the measurement period.² The majority of blackholing events are deactivated (strictly) through prefix withdrawal as opposed to through a re-announcement without a blackholing community tag. Specifically, we witness 1.294 M withdrawals, against 1.7 k re-announcements. Roughly 1.6 k (0.12%) of events are open-ended, i.e., are still active on the last day of our measurement period. We also find 6 k events that are deactivated both through withdrawal and re-announcement, which can occur if the event is inferred from BGP events on multiple collectors.

Figure 6.1 shows, per prefix length in the data set, the number of unique prefixes. A /32 is most-prevalent, with 137 k out of 146 k prefixes (93.7%).

¹Not all blackholing announcements propagate as far as public BGP collectors, meaning that we cannot possibly infer all blackholing events [42].

²The 2 collectors that did not provide us with any blackholing events are RV’s *KIXP* and *NAPAfrica*. The latter was added in February 2018 and thus only overlaps with our observation period for about a month. In fact, *RIPE NCC*’s *RIS* and *RouteViews* know a total of 43 collectors combined at current. *BGPStream* indexed 41 of them while we ran our analysis, of which we considered only 36 as 4 were not active during the studied period (*rrc02*, *rrc06*, *rrc08* and *rrc09*), and 1 is IPv6 only (*route-views6*).

Figure 6.1: The number of unique prefixes on the y-axis against the prefix length on the x-axis.

A /24 follows second with 8.2 k unique prefixes (5.6%). We note that the distribution of prefix lengths in our data set resembles that in Giotsas et al. [42], with the exception that /27s are less pronounced.

Figure 6.2: The number of blackholing events over time (black curves), the number of blackholed prefixes (gray curves), and origin ASNs (orange curves) of blackholed events in the data set.

Figure 6.2 shows blackholing event statistics over time. The *events* curve shows the number of events each day, with a daily average of 1184. The *unique prefixes* curve is noticeably lower the *events* curve, with a 384 daily average. This contrast stems from the fact that it is common for operators to (briefly) deactivate and reactivate blackholing to assess if DoS attack traffic is still being received [42]. This practice self-evidently raises the number of blackholing events. The average number of *origin ASNs* from which blackholing requests

originate daily is *100*. We will study blackholing events in more detail later in this chapter, once we have correlated blackholing activity with attacks data.

6.2.3 Blackholing Service Collateral

For our collateral damage analysis we need to map common Internet services (i.e., Web, mail and DNS) to blackholing events. Specifically, we need to map IP addresses of Web sites, mail exchangers, and authoritative name servers to blackholed prefixes. If we are to do this on the basis of DNS resource records, we require the **A** records of: (i) **www** labels³; and (ii) of (canonical) names found in mail exchanger (**MX**) and name server (**NS**) records. The OpenINTEL project, previously used a data source in Chapters 4 and 5, measures these records. As such, we use OpenINTEL-provided data.

We use data for all domain names under the three gTLDs: **.com**, **.net**, and **.org**. As we pointed out when we used these particular gTLDs previously, they combinedly account for roughly *50%* of all domain names in the global namespace [15]. Table 6.3 summarizes the OpenINTEL-provided data set (i.e., before mapping records to blackholed prefixes). The data set contains a total of nearly *230* million Web sites, as well as almost *41* M mail exchangers and *8.5* M name servers. We note that IP address to Web site mappings fully overlap with the DoS and blackholing data sets, but the coverage of **MX** and **NS** mappings is shorter as the functionality to resolve these records was added to OpenINTEL on January 22, 2017.

start	type	#names	#IPs
2015-03-01	Web	229.44 M	28.77 M
2017-01-22	Mail (MX)	40.92 M	5.17 M
	DNS (NS)	8.54 M	1.64 M

Table 6.3: DNS data for Web sites, mail exchangers and authoritative name servers. We observe a total of *229.44* M Web sites, and *40.92* M and *8.54* M mail exchangers and name servers.

6.2.4 Reactive Measurements

In addition to creating a large data set of blackholing events for “offline” analysis, we target a selection of blackholed prefixes with reactive measurements. Our measurement system enables us to do this in near real-time by using *BGPStream* in so-called *live* mode. In live mode, *BGPStream* yields new routing

³The existence of an **A** RR for the **www** label is taken as a Web site indicator.

data as soon as they become available through the *data broker*. This allows us to observe blackholing activity in near real-time, subject to an *observation delay*. The observation delay, which is five minutes at the very least, consists of the time it takes for a BGP collector to receive, store and make available the routing data, as well as the time it takes the *BGPStream* data broker to index the data and pass the BGP update to our measurement system.

Figure 6.3: A simplified view of the BGP activity of a single blackholing event, as well as the observation delays.

Figure 6.3 shows a simplified timeline of the activation and deactivation activity related to a single blackholing event. At time t_a , the requesting AS announces the blackholing request. δ_a is the observation delay that follows activation. Similarly, at t_d , deactivation is requested, and δ_d is the observation delay that follows. These delays show why we cannot launch active measurements in near real-time. Again, delays are five minutes at the very least, depending on the routing data project.

We launch additional measurements following both activation and deactivation. As we will explain later, this provides us with two points in time for which a comparison of results can either support or oppose the presumption that the observed blackholing event had any effect. For both types of observation, we launch two types of additional measurements in parallel. First, we send out traceroutes to each selected blackholed prefix, making use of the RIPE *Atlas* measurement infrastructure. This allows us to make inferences about the reachability of blackholed prefixes on the basis of traceroute measurements from multiple vantage points. Second, we attempt discovery of the services that we are interested in (i.e., Web, mail and DNS) by portscanning a handful of IANA-assigned ports. Because of the observation delays, we presume that the effects of blackholing – should there be any – will have kicked in, or worn off, before we observe the respective activity, provided that both activities are sufficiently spaced apart in time. For this reason we attempt to get measurement results as fast as possible (e.g., by requesting RIPE *Atlas* to start the measurement as soon as possible).

We apply several heuristics to determine if a reactive measurement should be launched:

1. We only reactively measure /32s. The reasons for this are that: (i) a /32 identifies a clear measurement target (i.e., the prefix covers exactly one IP address); and (ii) selecting only /32s helps us honor the measurement limits that RIPE *Atlas* imposes. RIPE caps concurrent measurements at 100, and expenditure at 1 M credits daily.
2. If activation and deactivation on the same prefix are not sufficiently spaced apart in time, we can expect results to overlap as a result of *Atlas* scheduling delays and BGP propagation delays. For this reason:
 - (a) We do not launch a reactive measurement after observing *deactivation* if activation on the same prefix was recently observed. This heuristic no longer applies once 90 seconds have passed – a waiting period that is commonly used to allow for BGP propagation.
 - (b) We do not launch a measurement upon observing *activation* if deactivation on the same prefix was recently observed. (In this case, we void the preceding reactive measurement too, because its results may be tainted by the effects of the now re-activated blackhole.) This also helps us respect Atlas limits, because we avoid creating successive measurements for a /32 in case operators rapidly deactivate and reactivate blackholing to see if an attack stopped – a known practice [42].
3. We do not launch more than four concurrent measurements for /32s that belong to the same /24. This too also helps us avoid Atlas limits in case many IPv4 addresses within the same /24 are blackholed successively.

We also void results if we determine that overlap has occurred after the fact. As an example, an activation measurement may return results later in time than a *deactivation* that is yet to be observed because of the delay while it has already occurred in real-time (cf. after t_d but before $t_d + \delta_d$ in Figure 6.3). This creates a form of overlap that we cannot prevent before launching the measurement. It is important to note that observation delays that we are dealing with (five minutes or longer) significantly reduce the number of blackholing events that can be reactively measured in a timely fashion. As our results will show, this cuts the number of usable measurements in about half (see Section 6.5). Having said that, it does not altogether prevent us from performing reactive measurements.

Probe Selection

Our system selects six *Atlas* probes to send traceroutes from. It attempts to select probes in distinct *peer*, *customer* and *provider* networks (two probes per relation). We renew the list of *Atlas* probes once daily, retrieving IPv4-capable probes that are stable for at least one day. We use CAIDA’s ASRank to determine network relations. In cases where not enough probes with a desired relation exist, we randomly select probes in neighboring networks. Completely random probes are selected as a last resort.

Service Discovery

In addition to launching traceroutes, we also scan a small number of ports in attempt to discover Internet-reachable Web servers, mail exchangers, and name servers. We probe the target /32 for the IANA-assigned ports of the services of interest. All three services have multiple ports or protocols assigned. We probe: 80/TCP and 443/TCP for Web; 53/UDP and 53/TCP for DNS; and 25/TCP and 587/TCP for mail.

Ethical Considerations

We send a small number of port probes and traceroutes per blackholing event. We expect half of the packets (i.e., those that follow blackholing activation) to not reach the target network. And in cases where they do, we expect them to not have an adverse effect given their limited numbers, even if paired with ongoing attack traffic. We therefore consider our reactive measurements to constitute slight harm at best for the target, and to not put volunteer hosts of *Atlas* probes at risk given the timing of the traceroutes sent from their connection.

6.2.5 Inferring Blackholing Efficacy

Inferring blackholing activity from BGP data suggests that blackholing is intended, but it does not guarantee that network traffic to the blackholed prefix is at all dropped. In other words, we cannot be certain that the blackholing provider honored the request to begin with. Moreover, it is possible that some parts of the Internet will be able to reach a blackholed prefix whereas others cannot. We use our reactive measurement data to support or oppose the presumption that the observed blackholing activity had an effect. We refer to this as (in)efficacy.

If we find a port to be in the open state exclusively after blackhole deactivation, we infer that the blackholing activity had an effect (i.e., we infer efficacy). If a port is found to be open upon activation, we infer inefficacy. Other cases are inconclusive, meaning that firewalled hosts, or hosts that do not run any

of the selected services publicly, will not enable us to make inferences. Please keep in mind that we perform portscans only from a single vantage point, which does not guarantee that the result holds for any source network. That is, our inferences do not necessarily support or oppose Internet-wide blackholing effects. Moreover, an ongoing attack may lead to port probes being dropped even though the blackhole was not effective – a hypothetical corner case that potentially skews the results.

Atlas traceroute results contain the `last-hop-responded` and `last-hop-is-destination` flags. A comparison of the combinations of these flags found on the activation and deactivation-related traceroute results allows us to support or oppose the presumption that the blackholing activity had an effect. If we find, from a given probe, that only the deactivation-related result sees `last-hop-is-destination`, we infer efficacy. If the activation-related measurement sees a `last-hop-is-destination`, we infer inefficacy. We do this per probe type to shed light on which source networks see an effect on network traffic, as well as regardless of network type.

6.3 Blackholed Attacks

We jointly analyze our data sets on attacks and blackholing to find “*blackholed attacks*”. In this analysis, we require an attack’s target IP address to be covered by the prefix of a blackholing event, and the attack’s start time to precede the blackholing event’s activation in time (by at most 24 hours). We will show later that blackholing is often triggered well within the hour following an attack’s start time.

Table 6.4 summarizes the matches. Surprisingly, we find more than 450 k attacks, towards almost 70 k targets (and involving 2.5 k ASNs) that were mitigated through blackholing. **This is the first large-scale empirical observation of DoS events and corresponding blackholing mitigation.**

source	#attack events	#targets	#ASNs
UCSD-NT	214.9 k (1.35%)	34.5 k (1.17%)	1732
AmpPot	241.0 k (1.97%)	47.5 k (0.79%)	2197
Combined	456.0 k (1.62%)	69.7 k (0.81%)	2543
Joint	18.4 k (4.12%)	5.7 k (3.25%)	800

Table 6.4: Blackholed Denial-of-Service attacks. This is the first large-scale empirical observation of DoS events and corresponding blackholing mitigation: 456 k of the 28.16 M attack events in our data sets are blackholed (1.62%), which involves 0.81% of all uniquely targeted IP addresses.

Only small percentages of the UCSD-NT and AmpPot data sets are blackholed, i.e., 1.35% and 1.97% of attacks, and 1.17% and 0.79% of unique targets. (Combined, we see blackholing for 0.81% of all unique target IPs.) While at first glance these small percentages might suggest that the data sets we examined contain “noise” (i.e., inferred attacks of negligible intensity), we show later in this section that even small intensities trigger blackholing. We thus conclude that such percentages reflect that: (i) we can observe blackholing only for a subset of ASes/targets; and (ii) blackholing adoption, while significant (2543 ASNs observed), might not be largely widespread. As future work we plan to further investigate this aspect, combining our data with blackholing at IXPs and the visibility of other community tags. Interestingly, for the 447.6 k attacks jointly launched against the same target (Table 6.1) that we observe in our DoS data sets, we find 18.4 k (4.12%) to be blackholed. This involves 3.25% (5.7 k) of unique target IPs, which, compared to 0.81%, leads us to believe that more serious attacks (i.e., those in which we observe the combination of multiple attack types) are more likely to be blackholed.

Our comparison of data sets also allows us to shed some light, for the first time, on the popularity of randomly-spoofed and reflection attacks compared to other DoS attacks (e.g., unspoofed) for which so far the research community has not been able to provide data on a global scale. Table 6.5 shows we find 159.9 k blackholing events preceded by a randomly spoofed attack, and 306.4 k preceded by a reflection attack. This means that we match 27.8% of all 1.30 M (Table 6.2) blackholing events in our data set with attacks. While, this preliminary result does not allow us to infer the fraction of different categories of attacks, it highlights that together **randomly-spoofed and reflection attacks represent a significant share of the attacks that operators dealt with in the last three years.**

attack source	#blackholing events	#prefixes
UCSD-NT	159.9 k (12.3%)	20.6 k (14.1%)
AmpPot	306.4 k (23.5%)	33.5 k (23.0%)
Combined	363.0 k (27.8%)	45.2 k (30.9%)

Table 6.5: Blackholing events that follow an (observed) Denial-of-Service attack in the UCSD-NT or AmpPot data sets, as well as for attacks in either. We match 363.0 k of 1.30 M blackholing events with attacks (27.8%).

Figure 6.4: Time until blackholing is activated. The distribution of the time between the start of attacks and the start of blackholing, for attacks in the UCSD-NT and AmpPot data sets. Almost half of all blackholed attacks (44.4%) see blackholing activated within a minute.

More than half of all blackholed attacks see mitigation activated within a matter of minutes

Figure 6.4 shows the time it takes for blackholing to be activated. For any blackholed attack in the UCSD-NT and AmpPot data, we analyze the delay between the start of the attack and the start of the associated blackholing event.⁴ For joint blackholed attacks – which may not see the randomly spoofed and the reflection attack start at the same time – we assume that the attack component that had started earlier in time triggered the blackholing event. To account for this assumption, we pick the longer mitigation delay for our analysis. In doing so we favor the risk of introducing “longer-than-actual” over “shorter-than-actual” times when estimating the delay with which blackholing starts. In other words, we pick an upper bound for the mitigation delay. It should be noted that we can only do this for joint attacks that we recognize as such, meaning that we cannot account for attack components that we do not observe (see Section 3.2.3). However, based on our observations of randomly spoofed attacks and reflection attacks, joint attacks are relatively rare.⁵

⁴BGP collectors, AmpPot instances, and the UCSD-NT infrastructure synchronize time through *NTP*. Notwithstanding, BGP timestamps are based on when the collector receives an update – not when the origin AS requested blackholing. Moreover, marginal time deviations may occur depending on where the BGP collector is in relation to the blackholing provider.

⁵We analyzed the start time differences between attack components of the 18.4 k joint blackholed attacks in our data (see Table 6.4) and find that 85.54% see the attack start spaced less than 40 minutes apart.

Nearly half of blackholed attacks (44.4%) see the blackhole activated within one minute, and 84.2% see activation within ten minutes. Such times suggest the use of automated detection and mitigation. Only for 0.02% of blackholed attacks it takes longer than six hours for blackholing to be activated.

Figure 6.5: The distribution of the time between the end of attacks in the AmpPot data set, and the end of correlated blackholing events. In 74.8% of blackholed reflection attacks, the blackholing is withdrawn in three hours or less after the attack stopped. In some cases, however, blackholing is left active for days after.

Often blackholing mitigation lasts way beyond the attack duration

Figure 6.5 shows the time between the end of blackholed attacks in the AmpPot data set and the end, i.e., deactivation time, of the associated blackholing event. (Blackholing “truncates” the attack end times in UCSD-NT data (see Section 3.2.1), which is why we do not analyze deactivation delays for randomly spoofed attacks.) We show that for 74.8% of blackholed attacks the blackhole is deactivated within three hours after the end of the attack. 96.1% of blackholed attacks see deactivation within 24 hours, meaning that for 3.9% it may take multiple days. These results suggest lack of automation in recovery from blackholing, and highlight that its side-effects (completely blocking *any* traffic reaching the victim) extend beyond the duration of the attack, i.e., a self-inflicted DoS. We will quantify this drawback at scale later in this chapter.

Figure 6.6: The intensity distribution for all attacks in the UCSD-NT data set (black curve), as well as for those that are blackholed (gray curve). We show that less intense randomly spoofed attacks are also mitigated – 13.1% see an inferred intensity of at most 3 Mbps (1 packet/s observed).

We see evidence that less intense attacks are also mitigated

Recall from Chapter 3 that the UCSD-NT data set contains a measure of attack intensity (pps_{max}), expressed in terms of the maximum number of backscatter packets per second observed. Figure 6.6 shows the overall distribution of intensities in the UCSD-NT data set, as well as the distribution for blackholed attacks only. 64.6% of blackholed attacks (gray curve) have an intensity not greater than 100 pps_{max} , which corresponds to an approximate attack traffic volume of 300 Mbps. This applies to 91.1% of all attacks (black curve), which confirms the intuition that attacks for which mitigation is observed are likely to be stronger. On that note, in Chapter 5 we already showed that stronger attacks lead to quicker outsourcing to protection services – the other global mitigation solution studied in this thesis. More importantly, however, a non-negligible percentage of blackholed attacks have a low intensity. Specifically, 13.1% see an intensity of at most 1 pps_{max} (3 Mbps). First, this result shows that operators mitigate – with such an extreme measure as blackholing – even less intense randomly spoofed attacks; which raises the question of what is the minimal effort needed by an attacker in order to induce the victim to recur to “shut down” an IP address for a certain period of time. In addition, this is the first time we are able to confirm (on a large scale) that even the smallest attack intensities inferred through a methodology based on indirect and partial observation of DoS phenomena but largely used in literature (Moore et al. [80]) are relevant, since they trigger mitigation. Finally, this result underpins the validity of the surprisingly large number of DoS attacks that we discovered in Chapter 3,

contributing to the bigger picture, and it provides a reference threshold to be used in the context of monitoring and situational awareness.

Figure 6.7: For the five most-used reflector protocols, the intensity distribution of all attacks in the AmpPot data set (upper plot), as well as for those that are blackholed (lower plot). We show that less intense reflection attacks are also mitigated. For example, 50% of all blackholed SSDP-based attacks see at most *4.9 requests/s*.

An analysis of the AmpPot data attack intensity measure (rps_{avg}) for blackholed reflection attacks yields similar results. The top five reflector protocols in the AmpPot data are: (1) NTP – 40.7%, (2) DNS – 25.6%, (3) CharGen – 22.6%, (4) SSDP – 8.3%, and (5) RIPv1 – 2.6%. We consider only these protocols and note that they are used in all but 0.2% of AmpPot attacks. Figure 6.7 shows the intensity per protocol for the top five reflection attack protocols for all AmpPot attacks as well as for those that are blackholed ((1) NTP – 45.0%, (2) DNS – 33.9%, (3) CharGen – 11.2%, (4) SSDP – 7.5%, and (5) RIPv1 – 2.1%). We here too show that operators also mitigate less intense reflection attacks (e.g., 4.9 rps_{avg} for fewer for 50% of blackholed SSDP-based reflection attacks). We also confirm the intuition that mitigated attacks are likely to be stronger on average. Specifically, between all AmpPot attacks and those blackholed, the median rates for SSDP, DNS and CharGen increase with 0.8, 6.6 and

11.5 rps_{avg} respectively. RIPv1 and NTP reflection see stronger increases, by 55.5 and 329.9 rps_{avg} , respectively.

Given that attacks of various intensities can be launched jointly against the same target, one could hypothesize that a less intense attack will only be mitigated by a target – with such an extreme measure as blackholing – if it is joined by a high-intensity attack. We analyzed the intensity components in the 18.4 k joint blackholed attacks in our data (Table 6.4). 9.82% of the joint randomly spoofed attacks have an intensity in the 25-th percentile (which corresponds to an intensity of up to 2.55 pps_{max}). About a fifth of these attacks, 20.54%, were joined with a reflection attack that falls in the 12.5-th percentile of its respective, i.e., protocol-specific intensity distribution (e.g., up to 13.2 rps_{avg} for NTP). 40.71%, 68.39% and 86.79% of the aforementioned randomly spoofed attacks were joined with reflection attacks that have an intensity in, respectively, the 25-th, 50-th or 75-th percentile. The presence of low-intensity combinations in joint blackholed attacks corroborates that less intense attacks are also mitigated with blackholing.

Figure 6.8: the attack duration distributions for all attacks (black curves) and blackholed attacks (gray curves) in the amppot data (upper plot) and the ucscnt data (lower plot). we find that for randomly spoofed attacks, the average duration drops, which, given the attack-inference methodology, is indicative that blackholing is effectively stopping (at least part) of victim-directed traffic.

The blackholing communities we observe reflect actual traffic filtering

Figure 6.8 shows the duration distributions of all attacks and of blackholed attacks, for the AmpPot data as well as the UCSD-NT data. For reflection attacks (upper plot), the duration of attacks goes up for those for which we observe blackholing, with *41.6%* of blackholed attacks lasting ten minutes or longer, against *29.2%* for all attacks. This confirms the intuition that mitigated attacks are more substantial also in terms of duration. For randomly spoofed attacks, however, *64.5%* of blackholed attacks last ten minutes or shorter, against *55.5%* of all attacks (lower plot). The duration thus decreases. This might seem counter-intuitive at first, but we note that an effective blackhole will drop all target-destined traffic, including the packets that trigger backscatter. Consequently, the attack end time observed through backscatter may not reflect the actual time at which the attack stopped. In fact, none of the blackholed attacks last longer than *3.2 h* in our data. On the other hand, the end time observed in a reflector honeypot does not necessarily change as the result of effective mitigation, because the honeypot can still receive spoofed requests, even in the event where the victim no longer receives any traffic. The asymmetric increase and decrease in duration thus confirms that the BGP communities we observe reflect actual blackholing activity.

6.4 Blackholed Services

Based on previous considerations on the actual temporary loss of use of the victim IP address, in some cases even beyond the attack duration, we study the impact blackholing may have on the availability of common Internet services (i.e., Web, mail and DNS). We therefore consider blackholing events that involve prefixes in which Web sites, mail exchangers, or (authoritative) name servers are hosted. To this end we match blackholed prefixes against the DNS data set described in Section 6.2.3.

type	#prefixes	#names associated		
		overall	no-alt	ratio
Web	13.7 k (9.40%)	782 k (0.34%)	670 k	85.6%
Mail (MX)	2247 (1.54%)	180 k (0.44%)	177 k	98.3%
DNS (NS)	1176 (0.80%)	10 k (0.12%)	10 k	98.4%

Table 6.6: Web sites, mail and name servers hosted in blackholed prefixes. For the relatively small percentages of name associations that we find, *85.6* to *98.4%* do not have an alternative, non-blackholed IP address.

Table 6.6 summarizes the results of joining the blackholing events with DNS data. We find that *13736* blackholed prefixes map to Web sites at the time of blackholing, meaning that *9.40%* of all *146* k uniquely blackholed prefixes host Web sites (see Table 6.2). Over the entire observation period, *782* k Web sites (see *overall* in Table 6.6) are associated with blackholed prefixes. That means *0.34%* of the *229* M Web sites inferred in total (see Table 6.3). We find *2247* prefixes with MX associations. This involves *180* k mail exchanger names (*0.44%* of all *40.92* M MX names in the DNS data). For authoritative name servers, *10318* NS names (*0.12%*) map to *1176* blackholed prefixes.

Services can be redundantly hosted, i.e., be hosted on multiple IP addresses (and thus havem multiple records in the DNS). Not all IP addresses tied to a service are necessarily blackholed. We investigate this by studying the presence of records that point to non-blackholed addresses. The *no-alt* column in Table 6.6 indicates the number of names that are found to **not** have a non-blackholed IPv4 address at the time of blackholing (i.e., lower is better). Respectively for Web, mail, and DNS, *85.6%*, *98.3%* and *98.4%* of the names found (see *ratio* in Table 6.6) do not have an alternative IP address at the time of blackholing.⁶ While we find relatively small percentages of associations with respect to the full DNS data set, the majority of associations that we find do not have a non-blackholed IP address. It follows that the associated services may become – and remain for extended time – unavailable when blackholing is left active and no IP address change takes place.

Figure 6.9: The number of Web sites associated with blackholed prefixes at the time of blackholing.

⁶OpenINTEL measures from a single vantage point, which means that some IP address records may be missing. We do not expect this to drastically change the ratios.

Figure 6.9 shows how blackholed prefixes relate to the number of Web sites that are hosted in their address space at the time of blackholing. Each bar represents a range of the number of Web sites per prefix, in consecutive powers of ten. The height of each bar represents the number of prefixes that fall within the bar's range. (Prefixes that are blackholed more than once can contribute to more than one bar, depending on cardinality changes.) 5991 of 13736 uniquely blackholed prefixes with Web site associations (43.6%) map to one Web site at the time of blackholing. 40.6% do so for two to ten Web sites. For about 6.4% of prefixes we find an association with more than a hundred Web sites. The maximum is 70.7 k in a single prefix. In other words, blackholed prefixes less often map to many Web sites. While this shows that on average blackholed prefixes do not affect too many Web sites (10 Web sites or fewer in 84.2% of prefixes), in rare cases tens of thousands of Web sites are involved.

Figure 6.10: The number of Web sites associated with blackholed prefixes over the three-year observation period, with an average of 1500 daily, and a maximum of 111 k.

6.4.1 Web sites

Figure 6.10 shows Web site associations over time with a 1614 daily average and maximum of 111 k. If we consider Web sites without an alternative, non-blackholed IP address, we find a daily average of 1318 and a maximum of 71 k. We find that Web sites without an alternative IP address map to only 35 prefixes on average, underpinning, as did Figure 6.9, that a few blackholing events by, e.g., hosters can potentially affect many Web sites. Note that these prefixes can still host other Web sites that do have an alternative address setup in the DNS.

6.4.2 Mail and DNS

When it comes to mail exchangers associated with blackholed prefixes, we find a daily average over time of *741* for overall, and *739* for mail exchangers without an alternative, non-blackholed address (i.e., “*no-alt*”). The maxima for the two categories are, respectively, *67568* and *64745*. For authoritative name servers, we find an overall daily average of *34.3*, a “*no-alt*” daily average of *33.8*, and maxima of *2224* and *2196* name servers blackholed, respectively. Our quantification shows that common Internet services can become altogether cutoff as a result of blackholing.

6.5 Blackholing Efficacy

Earlier in this chapter we showed that the blackholing communities that we observe reflect actual traffic filtering. This notion, considered together with our quantification from the preceding section, support that blackholing can cause non-negligible collateral damage (e.g., a maximum of *71 k* Web sites within the considered namespace cutoff from at least part of the Internet). In this section we analyze our reactive measurement data set in attempt to further support or oppose the presumption that the blackholing activity that we observe reflects traffic filtering.

category	#groups	
total	16054	
overlap	7677	47.82%
not deactivated	941	5.86%
unusable	72	0.45%
usable	7364	45.87%

Table 6.7: Traceroutes and portscans metadata. We measured *16054* blackholing events reactively. Accounting for results overlap and errors, *7364* usable measurement *groups* remain for further analysis.

Table 6.7 summarizes our one-month reactive measurements data set. The *total* number of blackholed prefixes measured is *16054*. Each measurement *group* contains traceroutes and portscans launched after observing blackholing activation, as well as, potentially, after deactivation-related traceroutes and portscans. This means that the number of measurements performed within each group is considerably higher. *47.82%* of groups are not usable because we determine that measurement overlap is likely to certain (see Section 6.2.4). *941* groups do not relate to a prefix for which the blackhole was deactivated during our measure-

ment period, and a small percentage (0.45%) of groups is not usable because at least one of its measurements saw an error (e.g., RIPE Atlas probe assignment issues). We are left with 7364 (45.87%) usable measurement groups.

responded	#service							
	Web		Mail		DNS		Combined	
$a \cup d$	2866		464		528		2993	
a	211	7.36%	41	8.84%	64	12.12%	231	7.72%
d	2855	99.62%	462	99.57%	524	99.24%	2981	99.60%
$a \cap d$	200	6.98%	39	8.41%	60	11.36%	224	7.48%
$a \setminus d$	11	0.38%	2	0.43%	4	0.76%	0	0.00%
$d \setminus a$	2655	92.64%	423	91.16%	464	87.88%	2769	92.52%

Table 6.8: Service discovery results for Web, mail and DNS – for activation (a) and deactivation (d) portscans. In the majority of cases we only see an open port response after deactivation (cf. 87.88% to 92.64% for $d \setminus a$), supporting that the blackhole was, at least partially, effective.

6.5.1 Portscan Inferences

Table 6.8 summarizes the reactive portscan results. The $a \cup d$ column shows per service the number of measurement groups in which we find at least one open port. 2993 of such groups are found combined, which is 41% of all 7364 usable groups (see Table 6.7). The $d \setminus a$ row shows cases in which strictly the deactivation portscan finds the port open. Combining the portscan results of the three common Internet services, we find 92.52% of portscan results to support blackholing efficacy. We infer inefficacy for the other 7.48% . Note that the percentages of groups where strictly the activation portscan sees an open port, i.e., $a \setminus d$, are near zero, which supports our methodology. On the basis of portscans we infer efficacy for 37.60% (2769) of 7364 usable groups, whereas for 3.04% (224 of 7364) we infer inefficacy.

6.5.2 Traceroute Inferences

Table 6.9 summarizes the reactive traceroutes results. The *groups* column shows, per probe type and for all types (see *combined*), the number of groups that have a pair of activation and deactivation-related results sent from a given probe and type. The *efficacy* column shows the number for which we infer efficacy. The *inefficacy* column is for inferred inefficacy.

For all probe network types combined, we see usable traceroute pairs for 6826 groups (92.69% of 7364). We infer efficacy for 2182 (31.97%) of them,

probe network	#groups	inference					
		efficacy		inefficacy		$\cap$	
Peer	5026	1464	29.13%	442	8.40%	51	1.01%
Provider	5403	1568	29.02%	315	5.83%	41	0.76%
Customer	2018	334	16.55%	164	8.13%	42	2.08%
Rand. neighbor	4332	1314	30.33%	385	8.89%	34	0.78%
Random	1511	563	37.26%	135	8.93%	9	0.60%
Combined	6826	2182	31.97%	664	9.73%	113	1.66%

Table 6.9: Reactive RIPE Atlas traceroute results and inferences for probes in peer, provider, customer and random (neighbor) networks.

and inefficacy for 664 (9.73%). In some cases we see result pairs that create a contradicting inferences (see the $\cap$ column), e.g., one *customer* network can see network traffic towards the blackholed prefix dropped, whereas another *customer* does not (recall, we attempt to select two *Atlas* probes per network type). We note that these cases do not occur often – only in 1.66% of groups, which suggests that the effects of blackholing are fairly consistent given a network’s relation to the network in which blackholing is deployed.

6.5.3 Joined Inferences

Table 6.10 summarizes how the efficacy inferences from reactive traceroutes (t) and portscans (p) relate to each other. Considering either portscans or traceroutes, or both combined (see $p \cup t$), we are able to infer efficacy for 3439 of 7364 groups (47%). 80.52% (of 3439) we could infer on the basis of portscans alone; 36.55% we infer exclusively on the basis of portscans. For traceroutes, the percentages are 63.45% and 19.48%, respectively. This shows that portscans contribute relatively more to our findings than traceroutes. Note that 3439 reflects the number of reactively measured blackholed prefixes for which we infer that network traffic is dropped at least partially. Our single portscan vantage point does not enable us to draw global conclusions and, as the traceroute results show, we also infer that, at times, from the viewpoint of some network relations, network traffic can still reach the blackholed prefix.

measurement	#groups
$p \cup t$	3439
p	2769 80.52%
t	2182 63.45%
$p \cap t$	1512 43.97%
$p \setminus t$	1257 36.55%
$t \setminus p$	670 19.48%

Table 6.10: Joined inferences of traceroute (**t**) and portscan (**p**) results, supporting efficacy for 3439 blackholed prefixes.

6.6 Corroborated Collateral Damage

We now quantify service collateral damage only for the blackholing events for which we, in the preceding section, inferred efficacy based on our data set of reactive measurements. We refer to this as corroborated collateral damage.

Table 6.11 shows how the blackholed prefixes for which we infer efficacy map to services. We find *Web* mappings for 734 prefixes, involving a total of 30916 Web sites (recall: this is for **com**, **net**, & **org**). Of these Web sites, 14669 do not have an alternative, non-blackholed IP address. We find 107 prefix mappings for *mail*, involving 3533 mail exchangers. Moreover, we find *DNS* mappings for 46 prefixes and 323 authoritative name servers. With the exception of the *overall* to *no-alt* ratio for Web, the ratios do not differ much from those in the broader mapping of services to blackholed prefixes (see Table 6.6). We cannot offer a definitive explanation for this contrast.

type	#prefixes	#corroborated names			#affected
		overall	no-alt	ratio	
Web	734	30916	14669	47.7%	14669
Mail	107	3533	3527	99.8%	522
DNS	46	323	318	98.5%	708

Table 6.11: Collateral damage of common Internet services hosted in blackholed prefixes for which we inferred efficacy.

The *affected* column shows the number of Web sites or domain names affected by service interruption. For the type *Web* this is simply the number of Web sites in **com**, **net**, and **org** without an alternative, non-blackholed IP address. These Web sites will be unreachable for the duration of the blackholing event – at least, for part of the Internet (recall, if we infer efficacy the effect is not necessarily

Internet-wide). For *mail* and *DNS*, the *affected* column shows the number of domain names that rely on these servers **exclusively**. Exclusively here means that domain names with alternative mail exchanger and authoritative name server records are not considered. In other words, 522 domain names rely exclusively on one or more blackholed mail exchangers – none of which may be reachable for the duration of the blackholing event(s). And 708 domain names rely on one (rarely more) authoritative name server that may not be reachable.⁷ Our inferences thus allow us to corroborate for only relatively small numbers of, e.g., Web sites, that they have become cutoff from (part of) the Internet.⁸

As final piece of the puzzle, we consider Web sites, mail exchangers and name servers that have no alternative, non-blackholed IP address at the time of blackholing, but change IP address after being blackholed (and prior to the blackhole’s deactivation). Note that this is exploratory as we cannot definitively say that these changes are due to blackholing and not the result of, e.g., DNS-based load-balancing that was already in place. Table 6.12 summarizes the results. We find that 3856 Web sites change to a non-blackholed IP address (26.28% of 14669). For mail exchangers and name servers, 257 (7.28% of 3527) and 6 (1.89% of 318) change, respectively. These latter two changes affect, respectively, 47 and 2 Web sites.⁹

type	#avoiding names		#affected
Web	3856	26.28%	3856
Mail (MX)	257	7.28%	47
DNS (NS)	6	1.89%	2

Table 6.12: Services changing from blackholed prefixes.

We quantified collateral damage in terms of the number of affected Web sites in three large gTLDs, but note that names in other TLDs may also be affected because our DNS data set only accounts for 50% of the global namespace. Similarly, servers other than those that we considered that rely on DNS resolution may require interaction with the blackholed authoritative name server and can therefore become collateral damage for the duration of the blackholing event. That is, provided that no IP changes take place – which we show to be uncommon.

⁷Mail Transfer Agents typically try to resend mail for days, meaning that a cutoff mail exchanger may only cause a delay. Moreover, DNS caching may help delay, for some clients, issues with name resolution.

⁸Note that the MX and NS records in the DNS data cover more than **com**, **net**, and **org**, which makes the *affected* numbers for mail and DNS seem disproportionate.

⁹Please note that, for reasons explained in a previous footnote, this relates to 37 (not 257) mail exchangers and 2 name servers used by domains in **com**, **net**, and **org**.

6.7 Related Work

Several approaches to DoS mitigation have been proposed in literature [71, 118]. However, it is only with our analysis in Chapter 5 that a first characterization of the adoption of DoS mitigation solutions was measured at scale. Giotsas et al. [42] present a comprehensive characterization of BGP blackholing activity on the Internet, based on public and private BGP data. The authors also correlate blackholing activity with a limited number of well-documented attacks. Dietzel et al. [37] study blackholing at a large IXP and quantify the effects of blackholing on network traffic rates. Chatzis et al. [30] emphasize that IXPs play a key role in deploying blackholing. Our contribution in this chapter instead focuses on analyzing BGP blackholing activity in combination with Internet-wide observations of DoS activity.

DoS attacks have been the subject of detailed studies [79]. Chapter 3 of this thesis provides the first measurement-based characterization of DoS attacks carried out on the basis of diverse and large-scale data sources. We also retraced the major steps in DoS analysis, from the seminal paper by Moore et al. [80]. While the work in Moore et al. focuses on the analysis of backscatter, Krämer et al. [65] and Thomas et al. [103] focus on DoS attacks as seen from a set of amplification honeypots. Santanna et al. [99] and Krupp et al. [67] analyze DoS attacks generated by Booters, while Wang et al. [112] analyze Botnet-related attacks. The focus of this chapter is not on DoS attacks per se, but on the relation between DoS attacks and blackholing.

In terms of blackholing drawbacks, Hinze et al. [45] study how BGP Flowspec, which allows for more fine-grained filtering than blackholing, has the potential to reduce collateral damage. Vasudevan et al. [110] present a simulation-based study of various DDoS mitigation strategies and look at network traffic collateral damage trade-offs. Pack et al. [89] perform an extensive experimental evaluation to understand the collateral damage of source address prefix filtering (SAPF). We instead focus on quantifying service collateral damage of Internet-wide blackholing activity at scale.

6.8 Concluding Remarks

In this chapter we studied BGP blackholing, the second mitigation solution considered in this thesis. We analyzed, on an Internet-wide scale, the co-occurrence of DoS attacks and BGP blackholing. By fusing data sources we were able to reveal insights about the operational deployment of blackholing as a DoS mitigation strategy. Based on our analysis, we argue that BGP blackholing defense mechanisms can react extremely fast, thus appear to be highly effective at protecting the *network* involved. However, some blackholing events last far longer

than the duration of the related attack, thus being very hard on the *services* and *systems* involved.

We therefore studied the major drawback of blackholing: its coarse-grained nature, leading to self-inflicted Denial-of-Service. Using three common Internet services as a measure – Web, mail and authoritative DNS – we quantified potential collateral damage. To corroborate that services are cut off (part of) the Internet, we reactively measure blackholed prefixes for a one-month period. This enabled us to further support or oppose, in very specific cases, that the blackholing activity had an effect.

The portscan part of our methodology relies on open ports to infer blackholing efficacy. It stands to reason that it is likely to find selected ports open on a host, provided that the host run one of the considered services publicly. However, we investigated only three services and thus a small selection of ports. Our results strongly suggest that a significant percentage of blackholed prefixes do not host Web, mail and DNS services, meaning that we can reactively measure only a limited number of blackholed prefixes with this approach. Applying portscans for an extended selection of services is possible, but the ethics (among others) of aggressive scanning would have to be evaluated. Moreover, our traceroute-based inferences require that last hops respond to traceroute packets, which is often not the case. Again, this limits our inferences, and alternative heuristics such as path analysis may allow for better inferrecess.

All in all, this chapter contributes to a better understanding of the whole DoS ecosystem: *(i)* it validates and reinforces some findings from previous chapters; and *(ii)* it adds other pieces to the puzzle of the bigger picture of DoS attacks (attacks, defenses, impact, etc.).

Exposure to Direct Attacks

In the previous two chapters we studied two global attack mitigation solutions: (i) DDoS Protection Services (Chapter 5); and (ii) BGP blackholing (Chapter 6). Our BGP blackholing study revealed questionable operational practices (e.g., long retention times) that, combined with how blackholing works (i.e., all traffic is dropped indiscriminately), creates a drawback. Of protection services we know that it is imperative that in DNS-based setups, customers keep their origin a secret.

In this chapter we study a major drawback of using DDoS Protection Services: the fact that attackers can in many cases altogether bypass protection as the result of so-called “origin exposure”.

We will identify a comprehensive set of “exposure vectors” and use them to quantify bypassability at scale, for the world’s most popular Web sites, and for leading commercial protection services that have already made appearances in preceding chapters. We will then join our resulting data set of exposed customers with the attacks data from previous chapters to see if exposed origins are subjected to DDoS attacks.

The study discussed in this chapter was previously published as the paper “Measuring Exposure in DDoS Protection Services” in Proceedings of the 13th International Conference on Network and Service Management (CNSM’17) [56].

7.1 Introduction

In Chapter 5 we revealed a trend in the adoption of DDoS protection services on the Internet. Despite a thriving market for commercial providers, the use of protection services knows various pitfalls, some of which arguably leave customers of cloud-based protection services with a false sense of security. For example, as we pointed out in Chapter 5, there is evidence that DPS protect their Web site but leave their DNS infrastructure unprotected. In this scenario, attackers could effect Denial-of-Service through a successful attack on the DNS rather than on the Web site. Even in cases where the DNS infrastructure is protected, either by a DPS or through another managed DNS provider, using a managed DNS service does not guarantee sufficient protection. This is showcased, among others, by the attack on Dyn in 2016, which caused service interruption for major Internet platforms, especially in North America [44]. A lesson learned by some is that the reliance on a single managed DNS service is ill-advised [24], which standards and existing literature on DNS stability also stipulate [27]. Another example of a potential issue is that protection services may be altogether bypassed by sophisticated attackers. In this chapter we will focus on this particular problem.

The problem

As explained in Chapter 2, *network traffic diversion* is the key mechanism that allows protection to be outsourced to a DPS. That is, traffic must be routed through the security infrastructure of a protection service to effect cloud-based scrubbing. Figure 7.1 (repeated from Section 2.3.1 for convenience) shows a schematic of how DNS-based network traffic diversion works.

For a DPS to be effective in a DNS-based setup, it is of the utmost importance that the origin of a service (e.g., a Web server’s actual IP address) is known only to the protection service [75]. Moreover, it is recommended that requests to the origin from any source other than the DPS proxy (or a set of proxies) are dropped [1]. This brings about various problems. First, hiding the origin is a form of security by obscurity: all clients of a service (both malicious and legitimate) are expected to use the DNS to resolve a domain name to an IP address in the DPS infrastructure (i.e., that of the reverse proxy). The origin IP address needs only to be accessed by the reverse proxy and is to remain obscured from attackers. If this obscurity fails, the DNS may be bypassed, and DoS attacks can be launched on the origin directly [83]. Second, setting up a firewall can be neglected (see Figure 7.1), leaving the origin wide open to direct attacks. In some cases, setting up a firewall is a complicated or even infeasible endeavor, especially if a DPS deploys a large number of reverse proxies [70].

Figure 7.1: Schematic of DNS-based network traffic diversion

Finally, even if a firewall is properly in place, the advantage brought by the DPS (e.g., better provisioned infrastructure) may be lost when attackers learn the origin. In other words, a sizable volumetric attack directly on the origin is likely to saturate the network, firewall or not.

In this chapter we focus on answering the following questions:

- What is the extent to which DPS-protected Web sites are susceptible to direct attacks on the origin?
- Can we find (empirical) evidence that exposed Web sites are involved in attacks?

Our work extends to earlier work by Vissers et al. [111] with the following three contributions. First, we analyze origin exposure in the DNS at a larger-than-ever-before scale. Second, we consider novel, not previously investigated DNS vectors through which the origin can be exposed. Third, we study if origins are subjected to attacks after protection is outsourced, and characterize the attacks in question.

7.2 Data Sources

In Chapter 5 we devised a methodology to infer DPS use from OpenINTEL DNS data. We now use this methodology to create a data set of DPS-protected Web sites specifically for this chapter. We have also worked with data on DoS attacks in previous chapters, which means we have data to study if exposed origins are involved in attacks. The remaining data that we need relates to origin exposure, which is altogether new to this chapter. This is data in which we first analyze the presence of exposure vectors in the DNS of DPS-protected Web sites, after which we actively probe to check if exposure is (still) actual.

We discuss the new data in more detail in the following two sections. Section 7.3 details our comprehensive set of exposure vectors. And in Section 7.4 we outline our methodology to identify and verify exposure among DPS-protected Web sites. In Section 7.5 we put all data sources together and describe the resulting data sets.

7.3 Exposure Vectors

In the introduction of this chapter we explained that if security by obscurity in DNS-based traffic diversion is broken, direct DoS attacks may be launched on the origin. In this section we explain various vectors through which the origin can be exposed. We describe two novel vectors (see Sections 7.3.4 and 7.3.6) in addition to various vectors that were previously identified in the literature [75, 83, 111].

7.3.1 Third-Level Domains (3LD)

The label `www` is typically used to give a canonical name to Web sites. Consider `www.example.com`, which is a third-level domain (3LD) in the zone of `example.com`. A zone can contain a number of labels, at various levels. Domain name administrators can configure labels such as `ftp` and `mail` to give canonical names to, respectively, *FTP* and *SMTP* servers. If other services run on the same IP address as the origin Web server, then the IP address of the origin may be exposed through the Resource Record (RR) on non-`www` labels.

This may seem like a trivial error to make, but given that not all protocols are interoperable with the reverse proxy mechanism on which DPS providers rely, it is not uncommon to have labels pointing directly to customer infrastructure.

Techniques to reveal other labels include: simply brute-forcing them, using zone enumeration [96], or by monitoring queries for them (e.g., using passive DNS [114]). In addition, certain DPS providers are associated with predictable labels to the origin. A first example are labels in DPS documentation that

people copy verbatim while configuring the DNS for their domain.¹ Another example are labels that are (automatically) added to for non-proxiable protocols (see Section 7.3.2).

7.3.2 Mail Exchanger (MX)

The **MX** record of a domain name specifies the location of the domain's mail server. The name in the **MX** RR typically resolves to an IP address that is running a *Simple Mail Transfer Protocol* (SMTP) server. Mail can be dealt with by a mail provider such as *Google Mail*, but domains can also host their own mail server. In the latter case, as DPS providers do not proxy mail ports, the **MX** RR will resolve to a customer IP address. If the mail server runs on the same IP address as the origin Web server, the **MX** RR exposes the origin.

If a domain hosts its own mail server, the **MX** RR can specify a name within the domain's own DNS zone. In other words, the **MX** of **examp.le** could point to **mail.examp.le**. This necessitates a 3LD label, as explained in Section 7.3.1.² An example is shown below. The **www** label resolves to the DPS proxy, whereas the **mail** label exposes the origin.

www.examp.le	IN	A	10.0.0.1
examp.le	IN	MX	mail.examp.le
mail.examp.le	IN	A	172.16.0.1

7.3.3 Sender Policy Framework (SPF)

The Sender Policy Framework (SPF) [64] allows domain name administrators to combat forged e-mails that appear to originate from a domain but are in fact sent by a rogue or compromised mail server with no relation to the domain at hand. This is done by publishing SPF information in the zone of a domain by means of a **TXT** record with an SPF value indicator (i.e., *v=spf1* ...). This record can specify, among others, the IP addresses from which e-mail can legitimately originate. If the domain's e-mail server runs on the same IP address as the Web server, the origin is leaked. A simple example is below, where the specified IP address exposes the origin (recall, 10.0.0.1 is the reverse proxy).

¹An example of this **origin-www** from Akamai's documentation.

²CloudFlare formerly automatically configured the **direct-connect** label to avoid potential conflicts with **MX** records. Nowadays, it will be set to **dc-<rand>.example.com**, which is harder to guess by brute-forcing 3LDs, but will still leak from **MX** records.

<code>www.example</code>	<code>IN</code>	<code>A</code>	<code>10.0.0.1</code>
<code>example</code>	<code>IN</code>	<code>TXT</code>	<code>"v=spf1 ip4:172.16.0.1 -all"</code>

7.3.4 Name Server (NS)

The NS record specifies the location of the name server that is authoritative for a domain. This name server could be in-bailiwick, meaning that the NS record for `example` could point to, e.g., `ns.example`. If the name server runs on the same IP address as the Web server, the origin is leaked. A self-explanatory example is shown next.

<code>www.example</code>	<code>IN</code>	<code>A</code>	<code>10.0.0.1</code>
<code>example</code>	<code>IN</code>	<code>NS</code>	<code>ns.example</code>
<code>ns.example</code>	<code>IN</code>	<code>A</code>	<code>172.16.0.1</code>

7.3.5 Conflicting Records

Web site administrators may, while configuring DNS for diversion, neglect to remove IP address records that point to the origin. This could lead, e.g., to two A RRs on the `www` label, one pointing to the origin, and the other to the proxy. A typical example is where the root of a domain (i.e., its apex, or `@`) still points to the origin, while the `www` label is properly set to the DPS proxy. A self-explanatory example is shown below. We refer to these cases as “conflicting records”.

<code>www.example</code>	<code>IN</code>	<code>A</code>	<code>10.0.0.1</code>
<code>www.example</code>	<code>IN</code>	<code>A</code>	<code>172.16.0.1</code>

7.3.6 IPv6 Address (AAAA)

Web servers may be dual-stacked, which means they can be reached over IPv6 as well as over IPv4. Not many DPS providers support IPv6, which means the AAAA record of a `www` label could expose the origin, even if IPv4 traffic is properly diverted to the DPS. In these cases the origin might be attacked using its IPv6 address.

7.3.7 IP Address History

The exposure vectors explained thus far relate to the way the zone of a Web site is setup at the moment a prospective attacker inspects the current state

of the DNS. While the current state may no longer expose a Web site’s origin, historic DNS data still can. For this reason, administrators are recommended to use a “clean” IP address for the origin, i.e., one that is not publicly known once they have configured DNS-based traffic diversion [102]. Using the old IP address breaks obscurity more easily.

7.3.8 Is an IP address a sufficient indication of exposure?

A leaked IP address is not per se an indication of exposure. It might be that an address no longer corresponds to the origin, or requests from anything but the reverse proxy are filtered. For this reason an extra step needs to be taken to make sure that the origin has been found. We address this in our methodology.

7.4 Methodology

The methodology by which we study the extent to which Web sites expose their origin IP address through the DNS contains various steps. Figure 7.2 shows the overall flow of information in our methodology. We start with a selection of Web sites that are protected by a DPS. For these Web sites we find potential origin exposure in the DNS using a longitudinal data set of active DNS measurements from OpenINTEL. We then take the potential origin IP addresses and scrape them for Web content, both by using the DNS and performing a “regular” request (i.e., connecting to the DPS proxy), and by bypassing the DNS (i.e., connecting directly). We then apply a custom DOM-tree comparison method to see if retrieved Web content is similar, i.e., the potential origin in fact corresponds to the origin. We further detail these steps in the following sections.

7.4.1 Long-Term, Always-On Protected Web Site Selection

Our analysis starts with a selection of long-term, always-on Web site customers for a selection of DPS providers. We look at always-on customers only because on-demand customers reveal the origin by design when traffic diversion is not active. Our selection targets Web sites that have been a customer for at least three months on the day that we perform our analysis. We take this selection as representative for the customers of a DPS that have a stake in hiding their origin. We select these Web sites by applying the methodology to infer DPS use that we introduced in Chapter 5.

Figure 7.2: Steps taken to analyze if Web sites expose their origin IP address through the DNS

7.4.2 Deriving Potential Origin IP Addresses

Given a set of always-on customers we pull relevant DNS resource records from OpenINTEL-provided DNS measurement data. If the DNS of a Web site is configured as such that the origin is leaked, the data set will reflect it.

Given the domain name `examp.1e`, we pull IP addresses (i.e., A & AAAA records) for the labels of interest: `www`, `@`, and for all the labels on the list of commonly-used labels mentioned in Section 7.5.1. We also pull IP addresses for the names in `examp.1e`'s MX and NS records. Moreover, we extract full IP addresses from TXT records that contain SPF information. We only consider full addresses in SPF information and not network blocks, e.g., a `/24`. Furthermore, SPF includes are not followed, because includes typically lead to third-party IP addresses.

We also pull historical IP addresses from the DNS data, on days that precede the day on which `www.examp.1e` became customer of the protection service. We filter all IP addresses that are invalid (e.g., `400.3.2.1`), within private address space (e.g., `192.168.0.0/16`), or those that are routed to an autonomous system (AS) number of a DPS provider.

7.4.3 Scraping Potential Origin Addresses for Content

We feed the set of always-on Web sites and potential origin addresses to a custom-built scraper, which sends “regular” requests to fetch content through the DPS reverse proxy. It also bypasses the DNS and sends direct requests to each and every IP address that potentially corresponds to the origin. Our scraper is built in Python, uses multi-threading, and has a requests pacing and backoff mechanism to avoid stressing hosts with repeated connection attempts or Web requests. Our scraper stores all results (e.g., HTTP status codes, connection errors, and content). Any content accompanied by an OK status code is fed to our origin verification system once scraping has completed. For other status codes we infer that the potential IP address does not correspond to the origin. We describe the origin verification methodology next.

7.4.4 Origin Verification

An OK status code on Web content does not imply that we found the origin. To verify if the origin was found, we make, for any given Web site, pairwise comparisons between the content that was returned by the DPS proxy and that of each of the requests that bypassed the proxy. As Web site content can vary with every page load (e.g., ads or dynamic content such as page generation timestamps) we cannot simply do a one-to-one comparison between the contents of two requests. Instead, we account for variable change by relying on a DOM-tree comparison, the motivation for which is that variable change in the content on a Web site modifies some, but not all of its DOM-tree structure.

We base our comparison on the the tree-edit distance algorithm by Zhang and Shasha [119], which counts the number of edit operations (inserts, deletes, and substitutions) to get from one tree to another. Rosiello et al. [97] used this algorithm to compare phishing Web sites to their original, and Vissers et al. [111] use it in a work similar to ours.

As an example of how the distance algorithm works in its pure form, consider the two trees in Figure 7.3. The labels of non-leaf nodes correspond to HTML or XML elements and the leaf labels correspond either to strings or empty elements. The tree’s parent-child relations correspond to nesting. Text leafs are formed by text on Web pages (e.g., in a link) or by comments in the HTML source (i.e., `<!-- text -->`). In Figure 7.3, the date stamps in the leaf labels are page generation times. The difference between the two trees is an added link (i.e., `href`) and a different page generation date stamp.

Using pure Zhang and Shasha, it takes three operations to get from the left-hand tree to the right: two additions for the `<a . . .>` node and its *click me* text child, and one substitution for the date stamp. This yields an edit distance of 3. This distance is larger than one would want, because the page generation time

Figure 7.3: Two simple DOM trees to compare for distance

labels are considered unequal, i.e., contribute +1 to the distance as a substitution. To address this issue we apply a string distance function to text labels. We use *Sift4* [19] for this purpose, which approximates the Levenshtein distance, but is faster. We consider attribute names (e.g., `style=...`) when we compare non-leaf nodes, but not the attribute values. Our rationale for this is that variable change in, e.g., form nonces should not add to the distance. Finally, we normalize the edit distance to be able to compare distances between pairs of trees with different sizes. Using normalization, the distance, D_n , between the trees in Figure 7.3, with 6 and 8 nodes respectively, becomes:

$$D_n(t_1, t_2) = \frac{2}{6 + 8} = \frac{1}{7} \quad (7.1)$$

We feed the resulting [0..1] distance to a binary classifier that considers content similar (i.e., the origin is exposed) or dissimilar, by applying a straightforward threshold comparison. We explain the threshold selection next.

7.4.5 Threshold Selection

We assume that the distance function, D_n , follows two distributions: one with $\mu_s = 0.0$ for similar DOM trees, and the other with $\mu_d = 1.0$ for dissimilar DOM trees. Figure 7.4 shows the resulting curves. All distances that we calculate form a sample of either curve. Although we cannot tell to which curve a calculated distance contributes, the curves intersect and create a minimum at $\mathfrak{t}$, which is a reasonable threshold to use in the binary classifier, i.e., $D_n \leq \mathfrak{t} \implies \text{similar}$.

Figure 7.4: Theoretical distance distributions for similar and dissimilar content

7.5 Data sets

As previously mentioned in Section 7.2, we work with some familiar data sources in this chapter. All of the longitudinal data sets that we extract cover a period of 16 months (January 22, 2016 – May 22 2017). The following sections provide more details.

7.5.1 Active DNS Measurements

To investigate origin exposure among Web sites that use a DPS we rely on OpenINTEL data. The set of resource records that OpenINTEL queries for includes any RR type tied to the exposure vectors outlined previously, in Section 7.3 (e.g., MX, NS, and TXT).

We use a subset of the data that OpenINTEL measures. Specifically, we use measurement data for domain names that belong to Web sites that are on the Alexa’s Top 1 million. The details of the data set are shown in Table 7.1. The total number of Web sites seen over the 16 months period is 3.3 M. While one might expect this number to be closer to 1 million, we point out that it is significantly due to changes in the Web site ranking, especially in the long tail. The *data points* column shows the total number of 7.1 G collected data points (e.g., A, CNAME, and NS records). The *size* column shows the size of the compressed measurement data using Parquet columnar storage [18], with a total of 205.0 GiB.

start	days	source	Web sites	data points	size
2016-02-22	486	Alexa	3.3M	7.1G	205.0GiB

Table 7.1: Active DNS data set for the Alexa Top 1 million

provider	Web sites
Akamai	30.6k
CloudFlare	357.2k
DOSarrest	3.2k
F5	5.0k
Incapsula	19.6k
Level 3	13.0k
Neustar	30.0k
Verisign	6.6k

Table 7.2: DPS use among Alexa top 1M Web sites for each of the 8 DPS providers that we consider.

The OpenINTEL platform does not target an extensive set of labels by default. That is, it does not brute-force hundreds of 3LDs for any domain it measures. As outlined in Section 7.3.1, however, 3LDs may expose the origin of a Web site. To analyze this exposure vector we instructed OpenINTEL to send out queries for commonly-used labels. (As we will show in Section 7.6, we targeted only domain names of interest to avoid burdening the DNS.) The set of labels that we used to this end was provided by SIDN, the .nl registry operator, which used their ENTRADA platform [116] to identify the 1000 most-frequently queried labels at authoritative name servers for the .nl zone.³ To avoid sending queries for labels that are unlikely to have IP address records we first removed labels from the top 1000 that are nonviable (e.g., `_dmarc` & `_domainkey`). Among the top-queried labels we find, e.g., `www-origin` and `direct-connect`, which are DPS-specific labels to bypass diversion, as outlined in Section 7.3. Other labels include `mail`, `smtp`, and `ftp`.

7.5.2 DDoS Protection Services

We analyze the extent to which Web sites that are long-term, always-on DPS users (i.e., continuously divert traffic) expose their origin to direct attacks. Using our methodology from Chapter 5, we create a data set that accounts for all Web sites on the Alexa Top one million. The DPS use data set covers the use of the by now familiar set of leading DPS providers, as long as they support DNS-based traffic diversion [47]. This means we consider eight providers: Akamai, CloudFlare, DOSarrest, F5 Networks, Incapsula, Level3, Neustar, and Verisign.

³Any query for a 3LD for which a `NOERROR` was returned was counted (once) towards the label's rank, over a one-month period.

We point out that in a BGP-based traffic diversion setup, the DPS announces the origin’s IP address space. Consequentially, the *origin* need not be obscured, as any traffic sent to the origin IP address will be routed towards the DPS by design. For this reason we do not consider BGP-based DPS use in this chapter.

Table 7.2 shows the details of the DPS use data set in terms of the total number of Web sites that we associate with each of the eight providers, over a period of 16 months. Later, in Section 7.6, we will extract long-term, always-on Web sites from this data set, in accordance with our methodology.

7.5.3 DoS Attack Events

start	days	#events	#targets	#/24s	#/16s	#ASNs
2016-01-22	486	7.95M	1.28M	0.41M	23083	20096

Table 7.3: DoS attack events data

The third data set contains 16 months worth of DoS attack events inferred at the UCSD Network Telescope [14]. Table 7.3 summarizes the data set of attack events. We observe a total of about 8 million attacks over the 16-month period, targeting 1.28 M unique IP addresses in 20 k distinct Autonomous System (AS) numbers.

7.6 Results

7.6.1 Long-Term, Always-On Customer Web Sites

By our methodology, our analysis starts with a determination of long-term DPS customers that outsource protection in an always-on manner. Table 7.4 shows the number of resulting Web sites for the Alexa Top 1 million. Note that for CloudFlare, we randomly sample 1:10 out of 41830 always-on customers given the large always-on customer base. For the eight DPS providers combined we find a total of 10884 long-term, always-on Web sites to study further.

7.6.2 Origin IP Address Candidates

We extract from the OpenINTEL data set, for each and every selected Web site, the set of IP addresses that are origin candidates. We then filter out IP addresses of various categories in accordance with our methodology: 3556 IP addresses are discarded since they are in private address space, 1162 addresses are associated with a DPS autonomous system, and 0 are considered invalid addresses. We

provider	Web sites
Akamai	2100
CloudFlare	4183
DOSarrest	245
F5	265
Incapsula	2854
Level 3	1173
Neustar	39
Verisign	25
total	10884

Table 7.4: Per DPS the numbers of long-term, always-on Web sites that we analyze further

are left with potential origin IP addresses for 9260 out of 10884 Web sites (85.08%). That means the other 1624 Web sites have no potential exposure in their DNS configuration from the get-go. Table 7.5 shows the number of Web sites potentially exposed per vector, and the number of IP addresses found for each of the exposure vectors. It is worth noting that the total number of addresses is an upper bound on the number of requests that the scraper needs to perform, since exposure vectors for a given Web site can overlap on an IP address. In case of such overlap, the IP address needs to be scraped “directly” just once for the given Web site.

exposure vector	Web sites (%)	#IP addresses
third-level domain (3LD)	7928 (85.62%)	753.6k
IP address history	3744 (40.43%)	22.3k
SPF	2396 (25.87%)	6.5k
conflicting record	2314 (24.99%)	3.1k
mail exchanger (MX)	2091 (22.58%)	2.8k
name server (NS)	401 (4.33%)	1.0k
IPv6	173 (1.87%)	0.2k
total	19047	789.5k

Table 7.5: Number of potentially exposed Web sites per exposure vector along with the total number of IP addresses

7.6.3 Scraping Results

The selection of Web sites along with their candidate origin IP addresses were fed to our scraper. For 9170 out of 9260 potentially exposed Web sites we got an answer for the regular HTTP request. That makes *99.03%*. The other *0.97%* led to a connection error of some form (e.g., connection, request, or read timeouts). For *96.4%* of 9170 Web sites we got a `HTTP STATUS OK`. Most other responses were `NOT FOUND`, i.e., not the origin (anymore), or `FORBIDDEN`.⁴ We also saw various status codes specific to DPS providers such as `523` (*CloudFlare: Origin is unreachable*) and `521` (*CloudFlare: Web server refused connection*). Self-evidently, as Web sites without regular content provide us nothing to compare to, our scraper will skip any direct requests specific to those Web sites.

Figure 7.5: All pairwise DOM-tree distance calculations

7.6.4 Exposed Web Site Origins

To select a D_N threshold for our binary classifier we calculate the distances for all DOM-tree pairs in the scraping data (see Section 7.4.5). Figure 7.5 shows the results. We find a minimum at threshold $t = 0.2$. Using this t , we verify that for *40.5%* of Web sites we found the origin IP address, in accordance with our methodology. This comes down to 4408 out of 10884 Web sites. Table 7.4) provides a breakdown per exposure vector and per DPS. We will discuss the results for each individual exposure vector in the following sections.

The largest exposure vector are third-level domains

27.95% of the analyzed Web sites expose their origin IP address on one of the commonly-used labels. This comes down to 3042 of 10884 Web sites. Table 7.7 shows the five most-common labels on which the origin is exposed. The top two

⁴A `FORBIDDEN` could be given by the origin to requests from anything but the reverse proxy, but given the lack of content we presume no exposure.

DPS	exposure vector							U
	3LD	History	SPF	Conflict	MX	NS	IPv6	
Akamai	38.19%	21.67%	0.81%	34.43%	0.19%	0.10%	0.33%	54.05%
CloudFlare	30.91%	13.20%	5.71%	6.36%	17.07%	0.26%	0.31%	46.04%
DOSarrest	15.10%	3.27%	1.22%	2.45%	0.82%	0.00%	0.41%	19.18%
F5	12.83%	7.17%	0.38%	2.26%	0.00%	0.00%	0.00%	15.85%
Incapsula	23.76%	13.70%	2.94%	4.38%	2.87%	1.02%	2.59%	34.76%
Level 3	16.37%	5.54%	0.60%	12.87%	0.26%	0.17%	0.09%	21.74%
Neustar	10.26%	0.00%	0.00%	12.82%	0.00%	0.00%	0.00%	17.95%
Verisign	8.00%	4.00%	0.00%	8.00%	0.00%	0.00%	0.00%	16.00%
all	27.95%	13.70%	11.80%	7.40%	3.22%	0.88%	0.40%	40.50%

Table 7.6: Per DPS and per exposure vector, the percentages of Web site customers that expose their origin. The **all** row shows results independent of DPS provider, and the U column is for the 7 exposure vectors combined.

labels (i.e., **direct** and **origin-www**) are inspired by DPS documentation for non-proxiable traffic. The **cpanel** label in the third position is specific to *cPanel*, a control panel for Web sites.⁵ The **webmail** label is typically used for Webmail software that is served by a Web server from a document root other than the primary Web site. If domain name administrators forget to divert traffic on this label then it might expose the origin. We checked a few of these cases by hand and encountered instances of Webmail software (e.g., Roundcube⁶).

label	exposes (%)
direct	418 (13.7%)
origin-www	392 (12.9%)
cpanel	387 (12.7%)
webmail	381 (12.5%)
dev	372 (12.2%)

Table 7.7: Common origin-exposing 3LD labels

Exposure through IP address history comes second

We can find the origin of *13.70%* of the analyzed Web sites based on historic IP address data. This comes down to 1491 of 10884 Web sites. It should be noted that for the always-on customers that started using a DPS before the first day

⁵<https://cpanel.com/>

⁶<https://roundcube.net/>

of the data set, we do not have any historic addresses to pull from the active DNS data. As such, *13.70%* is a lower bound.

Conflicting Records

11.80% of Web sites expose their origin through a conflicting DNS record. Within this exposure type, the majority of Web sites (*94.39%*) have an IP address on the domain root (i.e., the `@`). A small percentage (*5.61%*) have a conflicting IP address on the `www` label. In a handful of cases, Web sites expose the origin on both. Akamai sees most customers that have this type of origin exposure, with *34.34%* of Web sites.

Mail Exchangers

By the mail exchanger record, *7.40%* of Web sites expose their origin. Table 7.8a shows the three most-common labels in MX records that expose the origin of Web sites. The most commonly used label is `mail`, predominantly used by *36.4%* of all Web sites associated with this exposure type. We find a lot of labels of the form `dc-<rand>` in MX 3LDs. As outlined in Section 7.3, these labels can be traced to CloudFlare. Because these labels are unique and usually occur once, they create a long tail in the ranking of label occurrences. CloudFlare is also the DPS that sees the highest MX exposure, at *17.07%*.

label	exposes (%)
<code>mail</code>	292 (36.3%)
<code>mx</code>	4 (0.5%)
<code>mx1</code>	3 (0.3%)

(a) MX labels

label	exposes (%)
<code>ns1</code>	27 (61.4%)
<code>ns2</code>	16 (36.4%)
<code>ns</code>	1 (2.3%)

(b) NS labels

Table 7.8: Common labels in mail exchanger and name server records that expose the origin of Web sites

Sender Policy Framework

Of analyzed Web sites, *3.22%* expose their origin through the SPF information that is published in the DNS. We encountered only IPv4 addresses in SPF records that expose an origin. That is, from all IPv6 addresses in SPF information, not a single address exposes a dual-stacked origin.

IPv6 IP address exposure

As for IPv6 exposure, of Web sites that are dual-stacked (i.e., those that have an AAAA record as well as an A record), only *0.88%* expose the origin on the IPv6 address. Most Web sites that expose their origin over IPv6 are Incapsula customers. Specifically, this applies to *2.59%* of Incapsula customers.

Authoritative Name Server

Finally, *0.4%* of Web sites expose the origin through a name server record. That is, the authoritative name server for the domain name of the Web site runs on the origin Web server. Table 7.8b shows the three most-common labels in NS records that expose the origin. The largest exposure of this type among DPS providers is for Incapsula, where *1.00%* of the Web sites expose the origin through a NS record.

As for Table 7.6, it should be noted that for some DPS providers the number of always-on Web sites is low (i.e., of the order of tens to only a few hundred), which might make the breakdown less representative for the providers in question.

A Look at Verified Origin IP Addresses

We investigated the IP addresses on which Web sites are exposed and found several addresses that are shared by multiple (exposed) Web sites. That is, even though Web sites are individually exposed through one of the exposure vectors, they end up on a shared IP address. One reason for this is that Web sites share the same owner. As an example, we found adult content Web sites that were seemingly related from similar names and which share hosting. Another reason is that Web sites are placed at a party that provides hosting (e.g., Google or Amazon). We have seen one case where as many as 22 Web sites could be traced to the same IP address from Amazon. The common-most autonomous system numbers associated with IP addresses on which Web sites are exposed can be linked to: Amazon, OVH, UnifiedLayer, Hetzner, and DigitalOcean.

7.6.5 Attacked Web Sites

We matched the Web sites for which we verified origin exposure with our data set of attacks and find that the origin of 843 of 4408 Web sites were attacked after the Web site had started outsourcing protection to a DPS. This comes down to *19%* of all exposed Web sites. Given that multiple Web sites (or even

non Web services) can be hosted on the same IP address, we cannot definitively say if the Web site was the target of the attack.

As we explained when we first introduced the UCSD Network Telescope as a data source (see Section 3.2.1), the attacks data set has per attack event an intensity attribute, expressed in terms of maximum packets per second (**pps** average) that the victim backscatters to the UCSD network telescope. This value ranges up to 310 k for the data set considered in this chapter.

The top 1 percent of all attacks see a **pps** rate of 200 or higher, which equals an approximate attack network traffic volume of 615 Mbps to the victim or more! 205 of exposed Web sites see an attack of this strength on their origin, showing that exposed origins of supposedly DPS-protected Web sites are subject to strong attacks that are not diverted to the protection service.

7.7 Related Work

In 2013, McDonald [75] brought attention to the fact that content delivery security (e.g., network traffic diversion to DDoS Protection Services) can in some cases be bypassed, leaving Web sites vulnerable to attacks. Later in 2013, Nixon and Camejo [83] attracted even more attention to this fact. In 2015, Vissers et al. [111] performed the first study into origin exposure at scale, for five DPS providers. Our study extends the aforementioned work. First, we analyze origin exposure in the DNS at a larger-than-ever-before scale. We consider Web sites on Alexa's top 1 million, and cover eight of the leading DPS providers for which we have already shown an increasing trend in adoption in Chapter 5. Second, we identify a comprehensive set of vectors through which an origin can be exposed in the DNS. Among these vectors are novel, not previously investigated vectors. Third and final, we match exposed Web sites with attacks in UCSD-NT data – a source on DoS activity first introduced in Chapter 3. This allows us to study if the origin IP addresses of exposed Web sites are involved with attacks after the Web sites in questions start outsourcing protection to a DPS.

7.8 Concluding Remarks

In this chapter we studied the extent to which Web sites that outsource protection to DDoS Protection Services by means of DNS-based network traffic diversion expose their origin IP address through DNS misconfiguration. Our study covers Web sites on Alexa's Top 1 million list that are long-term customers of eight leading protection services. We evaluated previously known as well as novel exposure vectors to find the potential origin IP addresses of Web sites at scale. By using a DOM-tree comparison method we were able to verify that

40.5% of the studied Web sites expose a (responsive) origin, which means that the protection of these Web sites can be bypassed, and that these Web sites are vulnerable to direct DoS attacks. As a consequence, the use of a DPS might not be as safe as most Web site operators expect. Our results reaffirm previously known errors, and identify novel DNS misconfigurations, all of which operators should address to fortify their defenses.

We also matched the exposed origin IP addresses that we found with a large data set of DoS attacks. Our results show that the origin of 19% of Web sites see (high-intensity) attacks after protection was outsourced, which indicates that correct management and configuration are needed to ensure the efficacy of DPS use.

Conclusions

In this final chapter we draw conclusions from the research presented in earlier chapters. We draw our main conclusions first, with respect to the three goals that we set out with at the start of this thesis. We then revisit and discuss in detail each research question defined in Chapter 1. Finally, we provide prospects for future research that can build on the research presented in this thesis.

8.1 Main Conclusions

The upsurge of DoS attacks has left many questioning how to best deal with the DoS problem, ranging from individual operators to governments. How big is the problem? What exactly are we defending against? And how are mitigation solutions operated in practice? These questions are among many to understandably ask. At the start of this thesis we identified various challenges surrounding such questions, some of which are more *technical* in nature, and some of which are more of a *societal* nature.

Our goals, which we will revisit in a little bit, relate mostly to technical challenges. And while we thus set out to focus on technical challenges, it turned out that overcoming some of them puts *societal* contribution within reach. As an example, consider our contributions to the CPB risk report on cybersecurity and economics for the Netherlands Ministry of Justice and Security [88]. We feel that having successfully collaborated on this report further validates our

work and gives meaning to it beyond its scientific contributions. Our efforts in fusing and processing diverse data enabled us to inform policy makers, the research community and network operators through the report, various top-tier publications, and other methods of dissemination.

The absence of scientific reporting at scale on the topic of DoS attacks was a driving force behind our work. And even with diverse data in hand, processing it and studying attacks and mitigation at scale is not a straightforward matter. To lower this burden for others, and in effort to facilitate broader and independent research, we made available our data set of DoS attacks to the research community [3].

Much of our work would not have been possible without the preexisting data sources that we identified throughout this thesis. We believe that for the community to advance its understanding of the DoS problem further, data sharing needs to happen on a larger scale. The preexisting data used for this thesis, while accessible if you know where to look, are scattered. Exchange is done through agreements between researchers and not so much in a centralized fashion. This should be improved, for example through (non-commercial) data exchange facilities. The *Information Marketplace for Policy and Analysis of Cyber-risk and Trust (IMPACT)* platform (introduced in Chapter 1) sets an example [7] of how data sharing can be facilitated. This project is funded by the US government and we note that due to changes in policy, the IMPACT project was almost shut down in December 2018 as the result of lack of funding. The Internet Measurement Data Catalog (IMDC/DatCat) set another example of how sharing can be facilitated [101]. Its developers stipulated that data sharing among researchers should be facilitated. The funding for DatCat ended in 2014 and in 2018, after twelve years of operation, the service was discontinued. More DoS-specific, DDoSDB [5] is an effort to share important information about DoS attacks to help victims, the research community and security services. Its goal is, among others, to facilitate legal attribution and improve detection and mitigation. In the OpenINTEL project we also try to facilitate research and promote independent efforts by making available data publicly. These efforts were recognized by the Research Data Netherlands coalition, which awarded us the Dutch Data Prize 2018 [94] for making (quote) “an extraordinary contribution to science by making research data available for additional or new research.” We were also runner-up for the *ISOC.nl Internet Innovation Award 2019* [50].

We believe that efforts to make measurement data more open should intensify in the future. Moreover, we consider it sensible to design data exchange facilities in such a way that their continued existence does not depend on a single funding agency. This also means that if one data provider pulls out of a multi-lateral sharing agreement, the exchange does not necessarily become unusable for the other parties. In terms of data sources, we feel that jointly developed

and funded measurement infrastructure may create synergy between research and security communities, and may therefore help society as a whole in dealing with emerging and ongoing threats.

We will now draw our conclusions for each of the three goals that we set out with at the start of this thesis.

8.1.1 Goal 1

Three prominent technical challenges that we identified relate to: *(i)* the *availability of (diverse) data* on DoS; *(ii)* the *fusing and analyzing data at scale*; and *(iii)* concerns surrounding existing reports (e.g., commercial reports). Our desire to tackle these challenges gave rise to the first goal of this thesis:

Goal 1: *to study the DoS phenomenon on a global, Internet-wide scale, and to identify, join and validate – where applicable – existing data to broadly report on the the DoS problem*

This goal has two equally important objectives: *(i)* identifying viable data sources; and *(ii)* using the data to study the DoS phenomenon on a global scale.

In terms of data, identifying the right data to support work planned down the road is crucial. Data sources can be scattered and therefore acquiring diverse data can require some effort. In this thesis we progressively extend our research with more data. Our first steps can be found in Chapter 3, in which we established a novel approach to start characterizing the DoS phenomenon. The first large-scale data sources that we identified provide indicators of DoS activity on a global scale. We continued working towards the data objective of goal 1 in Chapter 4, where we fused DNS data with attacks data to study the impact of DoS attacks on the Web. In Chapter 5 we inferred DPS use, enabling us to study global DPS adoption. By jointly analyzing inferred DPS use with DoS attacks we were able to shed light on factors that influence DPS adoption following attacks (i.e., migration). In Chapter 6 we inferred blackholing activity from public BGP routing data and fused the resulting data set with attacks data. By doing so we were able to show, for the very first time at scale, how blackholing is used in practice following DoS attacks. We also used reactive measurements to study the effect of blackholing to an extent. In Chapter 7 we studied the bypassability of protection services, which involved creating altogether new data sets, as well as building on methodologies and data from previous chapters.

An unforeseen but welcome side effect of our successful efforts to systematically fuse and enrich diverse data is that it allowed us to validate two preexisting methodologies to infer attacks. Specifically, by linking inferred attacks to various types of mitigation (Chapters 5 & 6) we were able to validate the attack

inference itself in addition to some inferred attack attributes, thereby contributing back to the research efforts behind the preexisting data.

With fused and enriched data we were able to study the DoS phenomenon on a global scale (the second objective of goal 1). Using diverse attacks data we first characterized DoS attacks on the Internet (Chapter 3). Our findings include eye-opening statistics about the DoS problem, revealing for example an average of 30 k attacks daily, as well as the fact that about one-third of all /24 networks estimated to be active on the Internet had seen one attack at the least over a two-year period. Some of our findings were covered by US and NL media outlets such as *TheRegister* and *Tweakers* [31, 109]. Our results in Chapter 4 reveal the potential impact of attacks at scale: about two-thirds of all Web sites under the largest top-level domains can be associated with attacked infrastructure. In Chapter 5 we revealed a significant trend in global DPS adoption and also shed light on factors that are or are not determiners for migration following attacks. As part of our study of how blackholing is used following attacks (Chapter 6) we laid bare operational practices that may lead to users being cutoff from Internet services unnecessarily. We also quantified this disadvantage by considering DNS data, and used reactive measurements to corroborate effects in specific cases. In Chapter 7 we studied the bypassability of protection services. We showed that 41% of popular Web sites¹ that outsource protection are susceptible to attacks that bypass the DPS. We also confirmed that bypassing attacks occur on the Internet. This leads us to conclude that operators who outsource protection may have a false sense of security and not realize that sophisticated attackers are able to bring their Web sites down more easily than expected.

8.1.2 Goal 2

In the introduction of this thesis we pointed out that knowledge about the adoption of mitigation solutions on the Internet, as well as how they are operated and deployed when attacks occur, was missing. Our second goal was defined with this in mind:

Goal 2: *to study the adoption of DDoS Protection Services and BGP blackholing, and to investigate the operational practices of operators that use these solutions*

We started out with two expectations. First, we expected that the use of protection services would be on the rise. Second, we expected that blackholing, which is known to be coarse-grained, would only be used as a last resort to countermeasure attacks.

¹Popularity here means inclusion on Alexa's list of Top 1M most-popular Web sites.

Our first expectation was met. In Chapter 5 we studied the adoption of protection services on the Internet. We considered potential users (i.e., domain names) among 50% of all domain names in existence, for nine leading commercial providers. We showed a significant trend in adoption over time: a relative growth in DPS use of $1.24\times$ in contrast to a mere $1.09\times$ expansion of the considered namespace. We also revealed how operators use protection services, and found that while they protect Web sites, they leave the DNS infrastructure, at times, left unprotected.

Our second expectation turned out to be wrong. In Chapter 6 we showed that nearly half a million attacks were blackholed over a three-year period. And while we were able to confirm the intuition that blackholed attacks are typically stronger, we also found a large number of low-intensity attacks to have been blackholed. We also discovered that the countermeasure is often left active hours following the end of an attack. Blackholing will in such cases arguably needlessly bring about a self-inflicted DoS, meaning that hosts may be cutoff from (parts of) the Internet unnecessarily.

We also revealed that DPS adoption is largely driven by third parties such as hosting providers, at times involving millions of customers. Even though *societal* challenges relating to *privacy* and *digital independence* are not the focus of this thesis, we feel that our results may give rise to concern here. First, because individual Web site owners appear to not have control over when to adopt a protection service, nor which provider to use, nor where their network traffic is routed. And second, because they may not be aware that such decisions are made on their behalf.

8.1.3 Goal 3

The final technical challenge that we pointed out relates to mistakes in deployment and operation of mitigation solutions. We argued that undesirable side effects may lead to a false sense of security among operators. With this in mind, we defined goal 3 as:

***Goal 3:** to study problems surrounding the use of mitigation solutions that result from mistakes in use and bad operational practices, and to investigate whether or not attackers seize on these as an opportunity*

As we have already touched upon in our concluding comments for goal 1, we discovered, in Chapter 7, a significant percentage of popular Web sites for which the protection service could be bypassed. More in detail, we enumerated various ‘vectors’ through which exposed origins can be found by sophisticated attackers. We then quantified exposure on an Internet-wide scale. While intuitively, attacks that bypass protection occur on the Internet, our results also provide

evidence of this, for the first time to the best of our knowledge. Altogether, we conclude that protection services may lead to a false sense of security among users, while configuration mistakes leave users vulnerable to attacks that bypass protection.

In Chapter 6 we quantified the extent to which BGP blackholing may cutoff common Internet services. We showed that tens to hundreds of thousands of name servers, mail exchangers and Web sites may become cutoff from the Internet as a result of blackholing. We also used additional, reactive measurements to corroborate collateral damage in specific cases. Combined with our empirical evidence of operational practices, we conclude that some operators self-inflict a DoS unnecessarily – and for extended periods of time.

8.2 Research Questions Revisited

In the three sections that follow we discuss each of the research questions defined in Chapter 1. Each section contains the set of research questions that relate to a particular goal.

8.2.1 Research Questions for Goal 1

In this section we discuss each of the three research questions that relate to the first goal of this thesis. The first research question for goal 1 is:

***RQ 1:** Which data sources on DoS do we need in order to study the DoS phenomenon on a global scale? Are there existing data that we can work with, fuse or derive from? Or do we need to gather new measurements?*

We studied this research question in nearly every chapter of this thesis. Specifically, in Chapters 3 through 7. In Chapter 3 we identified two data sources that can be used to infer Denial-of-Service activity on an Internet-wide scale. The first data source involves backscatter traffic from randomly and uniformly spoofed DoS attacks to a network telescope (or darknet). We used data from the UCSD Network Telescope [17], which is a sizable, largely-unused /8 darknet that passively collects, among others, backscatter from randomly spoofed DoS attacks. We implemented the (pre-existing) detection and classification methodology by Moore et al. [80] to infer DoS activity from telescope data. The second data source relates to DoS activity logged by a set of globally placed AmpPot [65] instances. AmpPot is an open-source honeypot that aims to track reflection and amplification DoS attacks by mimicking reflectors. As the two data sources account for different types of DoS attacks, they complement each other. We augmented the attacks data with two additional data sets: (i) NetAcuity Edge data for IP geolocation; and (ii) Routeviews Prefix-to-AS mappings

data for BGP routing information. These metadata are necessary to, e.g., geo-locate targets or find the autonomous systems to which target IP addresses are routed. So as far as RQ 1 and Chapter 3 are concerned, we identified various data sources to work with. Some data sources required (significant) further processing, either by using pre-existing or new methodologies. In some cases the data were readily available.

To study the effect of attacks in Chapter 4 we added another data source to those used in Chapter 3 already. We took active DNS measurement data from the OpenINTEL project [20] to resolve the link between DoS attacks and Web sites. We fused this information with our attacks data so as to study the potential effect of attacks on the Web.

For Chapter 5 we needed data suitable to study the adoption of DDoS Protection Services on the Internet, as well as to investigate factors that influence migration. We devised a methodology to create a data set on the use of DDoS Protection services. This methodology built on data from the OpenINTEL project, along with BGP routing information.

Chapter 5 focused on one particular mitigation solution. We needed additional data to study the other in Chapter 6. To study BGP blackholing we used BGP data from two projects that make available such data publicly: (i) University of Oregon's RouteViews Project [16]; and (ii) RIPE NCC's Routing Information Service [11]. Both of these projects gather Internet routing data from globally dispersed collectors. We inferred a data set of blackholing events by looking for BGP announcements that are tagged with communities that are likely to signal blackholing. We combined our data set of blackholing events with the previously used data on DoS attacks to study BGP blackholing in terms of operational practices. For the part of Chapter 6 that relates to collateral damage of blackholing, we needed to create a data set based on reactive measurements, while taking the blackholing activity that we observed in public BGP routing data as a starting point. Moreover, for this study we also used data from the OpenINTEL project to map blackholed networks to Internet services (Web, mail and authoritative DNS).

In Chapter 7 we further worked with our data on the use of DDoS Protection services (Chapter 5), OpenINTEL data, and DoS attacks data. But to find potentially exposed origins of DDoS Protection Service users, and to verify whether we had found an exposed origin, we needed to create three additional data sets. First, we needed data on commonly-used third-level labels in domain names (e.g., the *3L* in *3L.example.org*).² Secondly, we needed to instruct the OpenINTEL platform to measure these labels. Third and finally, we needed to

²The set of 1000 labels that we used was provided by SIDN, the .nl registry operator.

scrape Web sites and IP addresses to verify whether we had found an exposed origin.

To sum up our conclusions as far as RQ 1 is concerned: we worked with diverse data throughout this thesis. Some data sources were raw and required (significant) processing prior to analysis, using pre-existing methodologies where possible. Some data sets were readily available for fusion (through collaboration with other research projects). And some data sets had to necessarily be created from scratch.

Our second research question, RQ 2, builds on the data-related results of RQ 1:

RQ 2: *What does the DoS landscape look like on a global scale in terms of attack occurrence and attack types?*

Our study of the DoS landscape in terms of attack occurrence and attack types is for the most part covered by Chapter 3. After fusing data on randomly and uniformly spoofed DoS attacks, data on reflection and amplification attacks, and IP geolocation and BGP routing metadata, we extracted macroscopic as well as detailed insights about Internet-wide DoS activity. Our study of the DoS landscape led to eye-opening statistics. We witnessed a total of almost 21 million DoS attacks over a two-year period, meaning that the average number of attacks each day was almost 30 k. These attacks involved 6.34 million unique target IP addresses, which belong to approximately 2.2 million /24 IPv4 networks. This is more than a third of /24 networks estimated to be active over the Internet, which shows the expanse of targets, and underpins (together with more than 30 k autonomous systems attacked) that many network operators are facing a DoS threat. We also discovered more than a hundred thousand instances in which both randomly spoofed attacks and reflection attacks were simultaneously launched against the same target.

We studied which network and transport layer protocols had commonly occurred in randomly spoofed attacks during the studied observation period. We also found which protocols had been abused most for reflection. Our results showed overlap with commercial reports for the same period. For example, we found various vendor reports that also reported on NTP as the most-used reflection protocol. But our results were also different in some respect, which we attributed to the customer-specific nature of the data used by vendors against our global, independent view.

In terms of attack duration, we revealed that most attacks are short-lived, with a median duration of approximately 5–8 minutes (depending on the attack type). Attacks of over an hour were not uncommon, but attacks that last longer than a day were very scarce. We also investigated the intensity of attacks and revealed that, of 30 k daily attacks, more than a thousand had seen an intensity

north of the medium intensity. In the case of randomly spoofed attacks, this means an inferred network traffic volume of *300 Mbps* to potentially hundreds of gigabits per second. This amounts to traffic volumes that are far from easy to deal with and can only be thwarted with a provision of network resources infeasible for “regular” network operators.

A small part of RQ 2 is covered by Chapter 6, in which we shed some light – for the first time on a global scale and in a vendor-independent manner – on the popularity of different attack types. By comparing blackholing activity with attacks we showed that, together, randomly spoofed and reflection attacks represent a significant share of the attacks that operators had to deal with during the observation period of our study.

Our third research question, RQ 3, relates to attack targets and the impact of attacks:

RQ 3: *Which targets are involved in DoS attacks? And what is the potential impact that attacks have on these targets?*

We made a comparison between the countries that attack targets geolocate to and Internet address space utilization by country. This allowed us to reveal (in Chapter 3) that the most-attacked countries ranking follows space usage patterns in principle.³ This underpins, together with the sheer amount of /24 IPv4 networks targeted, that the DoS problem is not a fringe problem that is limited to certain networks or specific countries.

In Chapter 3 we revealed that the network port numbers standardized to serve Web content had commonly been targeted by randomly and uniformly spoofed DoS attacks. Our knowledge of this prompted us look at the potential impact of attacks in terms of Web sites. We covered this in depth in Chapter 4. Our DNS measurement data provided us with 210 million Web site to IP address mappings, irrespective of attacks. By associating target IP addresses with Web sites, we inferred that approximately 9% of all IP address targets mapped to at least one Web site at the time of the attack. Moreover, we showed that the presence of Web-specific ports was even more pronounced in attacks that target Web infrastructure.

Over two years, we found 134 million Web sites involved with attacks. That is a almost two-thirds of the 210 M Web sites considered! In more than a hundred cases, a million Web sites or more mapped to a single IP address at the time it was attacked. We found an average 3% of Web sites attacked daily, and 1.3% (of 210 M) targeted by attacks of medium intensity or higher. While

³There were notable exceptions to this. For example, France ranked higher in terms of the target geolocation ranking than in Internet space utilization. We attributed this to the fact that a large hoster in France was attacked.

we could not determine if a single Web site or the hoster had been the intended target, our results underpinned that the number of Web sites potentially affected by attacks is considerable.

8.2.2 Research Questions for Goal 2

In this section we discuss the two research questions for the second goal of this thesis, which concerns mitigation solutions. The first research question relates to DDoS Protection Services:

***RQ 4:** Can we quantify the adoption of commercial, cloud-based DDoS Protection Services? In which manner do customers use such services? And what are the factors that drive adoption?*

We tackled RQ 4 in Chapter 5 of this thesis, in which we presented a methodology to infer the use of DDoS Protection Services from DNS and BGP measurement data. In our study we considered nine leading commercial providers. Our analysis revealed a significant trend in adoption over time. Specifically, under the three largest gTLDs (i.e., `.com`, `.net` and `.org` – representative for about half of all domain names in the global namespace), we showed a relative growth in use of $1.24\times$ against an overall domain name growth of $1.09\times$. We also revealed upward trends for domain names under the country-code TLD `.nl` and domain names on Alexa’s Top 1M list.

We also showed that adoption of protection services is largely driven by third parties such as Web hosting providers. Some of the examples that we highlighted involve millions of domain names. This finding is particularly relevant to RQ 6 so we will discuss it further later.

Later in Chapter 5 we presented a taxonomy to describe the causal relation between the adoption of protection services and DoS attacks. Our taxonomy helped us quantify the extent to which prospective users adopt (i.e., migrate to) a protection service after being involved with an attack. We showed that migration is more likely following attacks and we highlighted that high-intensity attacks create an urgency to migrate quicker – 80.7% of victims of very strong attacks migrated within a single day.

In terms of the manner in which customers use protection services, we showed how DDoS Protection Services are used in practice. We broke down various forms of setting up network traffic diversion (a requirement for use) and also quantified the dynamic nature of protection at scale, i.e., whether protection is outsourced in an *always-on* or *on-demand* manner. To our surprise we discovered that some users tend to protect their Web site but not their DNS infrastructure – a finding that we will revisit with the discussion of RQ 6.

Our next research question relates to the second mitigation technique that we studied, BGP blackholing:

RQ 5: *How widespread is the use of BGP Blackholing for the purposes of DoS mitigation? And how do users, i.e., network operators, use blackholing when faced with DoS attacks?*

In Chapter 6 we studied how often (and how) blackholing is used as an operational countermeasure to mitigate DoS attacks. We started by inferring from public BGP routing data a data set of blackholing events. This resulted in a sizable data set of 1.3 million blackholing events, covering a three-year period. We matched the blackholing events with DoS attacks from our (previously identified) complementary data sources on DoS activity and discovered only a small percentage of attacks mitigated through blackholing (1.62% of 28.14 M attacks). While such a small percentage suggests that blackholing is not prominently used to mitigate attacks on the Internet, we did find 2543 autonomous systems involved, showing significant adoption of blackholing among operators.

We also shed light on various operational habits surrounding blackholing. We found that the countermeasure is typically put in place rapidly, with 44.4% of attacks blackholed within a single minute. This showed us that defenses are automated and capable of responding quickly, which is very positive. However, in terms of recovery, we discovered that the countermeasure is often left active beyond the end of the triggering attack. Surprisingly, in 3.9% of cases this was even done for over 24 hours. What this showed us is that some operators choose to let the side-effects of blackholing (completely dropping *all* traffic) extend well beyond the duration of an attack. We also found evidence that less intense attacks are blackholed. In fact, in 13.1% of cases we witnessed an (inferred) network traffic intensity of at most 3 Mbps, which raises the question how much effort an attacker has to do to set in motion a Denial-of-Service self-inflicted through blackholing.

8.2.3 Research Questions for Goal 3

In this section we discuss the two research questions for the third and final goal of this thesis, which concerns problems with mitigation. The first research question relates to problems surrounding the use of DDoS Protection Services:

RQ 6: *Can we identify problems with the adoption of DDoS Protection Services? Can we quantify this problem on the Internet? And do we see evidence that attackers actively seize on potential problems?*

In Chapter 5 we discovered that some DPS users do not protect their DNS infrastructure, which in essence provides attackers with an alternative attack

surface to attempt to achieve Denial-of-Service. This underpins concern as to whether users who outsource protection are sufficiently protected. Our finding prompted us to investigate other means by which a DPS can be bypassed.

In Chapter 7 we studied the extent to which the protection of popular ⁴ Web sites can be bypassed. We took vectors by which the origin IP address of a supposedly protected domain name was known to potentially leak through DNS (mis)configuration. We complemented these vectors with novel DNS exposure vectors and then showed – at scale – that 41% of popular Web sites had exposed their origin, leaving them potentially susceptible to direct DoS attacks. We went a step further and combined our knowledge of exposed origins with data on DoS activity. In doing so were able to show that origin IP addresses are indeed targeted with DoS attacks.

In Chapter 5 we showed that large parties such as hosting providers drive the adoption of protection services. It appears to then be the hosting provider that decides when to outsource protection. Control over where network traffic is diverted may not lie with the Web site owner. In fact, the Web site owner may not even be aware that such decisions are made on their behalf. We note that this could be a privacy concern for reasons outlined in Section 1.3.6.

Our next and final research question is about problems with BGP blackholing:

***RQ 7:** Can we quantify the adverse effects of blackholing on the Internet?*

In the first part of Chapter 6 we unveiled two potentially problematic practices surrounding the use of blackholing. First, we showed that the counter-measure is oftentimes left active for many hours following an attack. Second, we showed that it is common to blackhole very low-intensity DoS attacks. Both these operational practices raise concern as to services being cutoff from the Internet unnecessarily.

In the second part of Chapter 6 we therefore quantified the extent to which BGP blackholing may cutoff Internet services, which we refer to as “service collateral”. We focused on three common Internet services: Web, mail and authoritative DNS.

For a three-year period, we found that 670 k Web sites were potentially cutoff from the Internet during blackholing activity. For mail exchangers and authoritative name servers we found 177 k and 10 k names, respectively. These results highlight the scale at which Internet services may become cutoff as a result of blackholing.

We also presented a methodology that uses reactive measurements to corroborate, in specific cases, that observed blackholing activity has, to some extent,

⁴Popularity here means inclusion on Alexa’s list of Top 1M most-popular Web sites.

an effect on network traffic. Our methodology is only applicable in specific cases because it requires hosts to respond to traceroute packets or probes to a conservative selection of ports, which we found not to be common in our data. Notwithstanding, our inferences more often supported than opposed blackholing efficacy.

8.3 Prospects for Future Research

In this last section of our final chapter we discuss prospects for future research. We imagine several directions to advance our understanding of the DoS problem and to increase situational awareness about threats to Internet stability and reliability.

- We provide a comprehensive view of randomly spoofed and reflection and amplification attacks, yet we found corresponding attacks for only 28% of blackholing events (see Section 6.3). These results suggest that, even equipped with two large-scale, complementary data sources that provide indicators of attack activity on the Internet, we still do not see the total extent of the DoS problem. We find this particularly striking and consider the development and systematic integration of other attack data sources (e.g., unspoofed volumetric attacks and semantic attacks) as a direction for future research. We hope that by demonstrating the utility of our data fusion approach in this thesis we can inspire the community to consider what would be required to expand the set of data sources to further advance our understanding of the DoS problem.
- This thesis would not have been possible without preexisting data sources. For our analyses we extracted data sets with longitudinal yet delimited observation periods. This means that most of our work is subject to a bounding problem. As an example, consider that we do not know which attacks took place earlier in time, or which mitigation was triggered later in time. Fusing data on a continuous basis would allow for this problem to be eliminated and also provide opportunity for sustained situational awareness. We therefore consider continuous, systematic data integration as a future work directive.
- In terms of data sharing, we consider the further development of non-commercial data exchange facilities as a much-needed development. We do, however, believe that prudence is in order. First, we believe that data sharing facilities should be structured in such a way that their continued existence does not depend on a single party. In other words, if one funding

agency pulls out, the exchange should not altogether collapse. And second, we feel that facilities could benefit from being designed with a *quid pro quo* mindset, meaning that there is a deterrent for any one data provider to withdraw.

Bibliography

- [1] “A Complete DDoS Protection Solution From a Leading Provider of Internet Infrastructure,” https://www.verisign.com/en_US/security-services/ddos-protection/index.xhtml, accessed: 2017-05-24.
- [2] “Arbor solutions,” <https://www.netscout.com/arbor>.
- [3] “CAIDA UCSD Randomly and Uniformly Spoofed Denial-of-Service (RSDoS) Metadata dataset,” https://www.impactcybertrust.org/dataset_view?idDataset=915.
- [4] “DDoS Prevention Services: Multi Layered DDoS Security Solutions,” <https://www.radware.com/solutions/security/>.
- [5] “DDoSDB: Collecting and Sharing the most important information of DDoS attacks,” <https://ddosdb.org>.
- [6] “Implementing BGP Flowspec,” https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-2/routing/configuration/guide/b_routing_cg52xasr9k/b_routing_cg52xasr9k_chapter_011.html, accessed: 2019-06-01.
- [7] “INFORMATION MARKETPLACE FOR POLICY AND ANALYSIS OF CYBER-RISK & TRUST,” <https://www.impactcybertrust.org/>.
- [8] “Lost in Space: Supplemental: Country Inequality (Interactive),” http://www.caida.org/publications/papers/2016/lost_in_space/supplemental/country_inequality/.
- [9] “MADDVIPR: Mapping DNS DDoS Vulnerabilities to Improve Protection and Prevention,” http://www.caida.org/publications/presentations/2019/maddvipr_dhsst/.
- [10] “Remotely triggered black hole filtering - destination based and source based (White paper),” https://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf, accessed: 2018-06-01.
- [11] “RIPE NCC Routing Information Service (RIS),” <http://https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>.
- [12] “RIPE78,” <https://ripe78.ripe.net/>.

- [13] “Routeviews Prefix to AS mappings Dataset (pfx2as) for IPv4 and IPv6,” <http://www.caida.org/data/routing/routeviews-prefix2as.xml>.
- [14] “The CAIDA UCSD Near-Real-Time Network Telescope – 2016-01 – 2017-05,” http://www.caida.org/data/passive/telescope-near-real-time_dataset.xml.
- [15] “The Domain Name Industry Brief,” https://www.verisign.com/en_US/innovation/dnib/index.xhtml, accessed: 2019-06-01.
- [16] “University of Oregon Route Views Project,” <http://www.routeviews.org>.
- [17] “UCSD Network Telescope (UCSD-NT),” 2010, http://www.caida.org/projects/network_telescope/.
- [18] “Apache Parquet,” 2014, <http://parquet.io/>.
- [19] “Super Fast and Accurate string distance algorithm: Sift4,” 2014, <https://siderite.blogspot.com/2014/11/super-fast-and-accurate-string-distance.html>.
- [20] “OpenINTEL Active DNS Measurement Project,” 2015, <http://www.openintel.nl/>.
- [21] “DDoS Threat Landscape Report 2015-2016,” Imperva, Inc., August 2016.
- [22] “DNS and Internet Naming Research Directions (DINR),” 2016, <https://ant.isi.edu/events/dinr2016/>.
- [23] “Global DDoS Threat Landscape Q4 2017,” Imperva, Inc., March 2018.
- [24] R. Abhishta Abhishta, van Rijswijk-Deij and L. J. Nieuwenhuis, “Measuring the Impact of a Successful DDoS Attack on the Customer Behaviour of Managed DNS Service Providers,” in *Proc. of the 2018 Workshop on Traffic Measurements for Cybersecurity (WTMC’18)*. ACM, 2018, pp. 1–7.
- [25] M. Abliz, “Internet Denial of Service Attacks and Defense Mechanisms,” Tech. Rep. TR-11-178, March 2011.
- [26] E. Addley and J. Halliday, “WikiLeaks supporters disrupt Visa and MasterCard sites in ‘Operation Payback’,” <http://www.theguardian.com/world/2010/dec/08/wikileaks-visa-mastercard-operation-payback/>, accessed: 2019-06-01.
- [27] M. Allman, “Comments On DNS Robustness,” in *Proc. of the 2018 Internet Measurement Conference (IMC’18)*, 2018, pp. 84–90.
- [28] D. Anstee, P. Bowen, C. Chui, and G. Sockrider, “Worldwide Infrastructure Security Report,” Arbor Networks, Inc., 2016.
- [29] R. Chandra, P. Traina, and T. Li, “BGP Communities Attribute,” RFC 1997 (Internet Standard), Internet Engineering Task Force, August 1996. (link)

- [30] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is More to IXPs Than Meets the Eye," *SIGCOMM Computer Communications Review (CCR)*, vol. 43, no. 5, 2013.
- [31] R. Cirgwin, "DoS scum attacked one-third of the 'net between 2015 and 2017," https://www.theregister.co.uk/2017/11/05/caida_study_finds_one_third_of_the_internet_suffered_denial_of_service_attacks_between_2015_and_2017/, accessed: 2019-07-01.
- [32] k. Claffy and D. Clark, "The 9th workshop on active internet measurements (AIMS-10) report," *ACM SIGCOMM Computer Communication Review*, vol. 47, no. 5, pp. 35–38, 2017.
- [33] k. Claffy and D. Clark, "The 10th workshop on active internet measurements (AIMS-10) report," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 5, pp. 41–47, 2018.
- [34] J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir, "Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks," in *Proceedings of the 2014 ACM Internet Measurement Conference (IMC'14)*, 2014, pp. 435–448.
- [35] A. Dainotti, K. Benson, A. King, B. Huffaker, E. Glatz, X. Dimitropoulos, P. Richter, A. Finamore, and A. C. Snoeren, "Lost in Space: Improving Inference of IPv4 Address Space Utilization," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 34, no. 6, pp. 1862–1876, 2016.
- [36] D. E. Denning, "Activism, hacktivism, and cyberterrorism: The Internet as a tool for influencing foreign policy," *Networks and netwars: The future of terror, crime, and militancy*, vol. 239, p. 288, 2001.
- [37] C. Dietzel, A. Feldmann, and T. King, "Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild," in *Proceedings of the 17th International Conference on Passive and Active Measurement (PAM'16)*, 2016, pp. 319–332.
- [38] B. Donnet and O. Bonaventure, "On BGP Communities," *SIGCOMM Computer Communications Review (CCR)*, vol. 38, no. 2, 2008.
- [39] D. Element, "Netacuity edge premium edition," <http://www.digitalelement.com/solutions/netacuity-edge-premium>.
- [40] F5 Networks, Inc., "2016 DDoS Attack Trends," November 2016.
- [41] A. Garg, J. Curtis, and H. Halper, "Quantifying the financial impact of IT security breaches," *Information Management & Computer Security*, vol. 11, no. 2, pp. 74–83, 2003.

- [42] V. Giotsas, P. Richter, G. Smaragdakis, A. Feldmann, C. Dietzel, and A. Berger, “Inferring BGP Blackholing Activity in the Internet,” in *Proc. of the 2017 Internet Measurement Conference (IMC’17)*, 2017, pp. 1–14.
- [43] A. Greenberg, “WikiLeaks Supporters Aim Cyberattacks At PayPal,” <http://www.forbes.com/sites/andygreenberg/2010/12/06/wikileaks-supporters-aim-cyberattacks-at-paypal/>, accessed: 2019-06-01.
- [44] S. Hilton, “Dyn Analysis Summary Of Friday October 21 Attack,” <http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/>, October 2016.
- [45] N. Hinze, M. Nawrocki, M. Jonker, A. Dainotti, T. C. Schmidt, and M. Wählisch, “On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP,” in *Proc. of the ACM SIGCOMM 2018 Conference on Posters and Demo (SIGCOMM’18)*, 2018.
- [46] R. Hofstede, M. Jonker, A. Sperotto, and A. Pras, “Flow-Based Web Application Brute-Force Attack and Compromise Detection,” *Journal of Network and Systems Management*, vol. 25, no. 4, pp. 735–758, 2017.
- [47] R. Holland and E. Ferrara, “The Forrester Wave™: DDoS Services Providers (Q3 2015),” Forrester Research, Inc., July 2015.
- [48] C. Huang, A. Wang, J. Li, and K. W. Ross, “Measuring and Evaluating Large-Scale CDNs,” in *Microsoft Research Technical Report MSR-TR-2008-106*, October 2008, (full paper withdrawn from the 8th ACM SIGCOMM Conference on Internet Measurement (IMC’08)).
- [49] H. Hueck, “Zorgen over cyberveiligheid van grote banken,” <https://fd.nl/ondernemen/1273770/zorgen-over-ddos-beveiliging-van-grote-banken>, accessed: 2018-10-16.
- [50] ISOC.nl, “ISOC.nl Internet Innovatie Award 2019,” <https://awards.isoc.nl/innovatie/2019/>, 2019.
- [51] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti, “Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem,” in *Proc. of the 2017 Internet Measurement Conference (IMC’17)*, 2017, pp. 100–113.
- [52] M. Jonker, A. Pras, A. Dainotti, and A. Sperotto, “A First Joint Look at DoS Attacks and BGP Blackholing in the Wild,” in *Proc. of the 2018 Internet Measurement Conference (IMC’18)*, 2018, pp. 457–463.
- [53] M. Jonker, “The drawbacks of blackholing,” 2019, <https://blog.apnic.net/2019/07/16/the-drawbacks-of-blackholing/>.

- [54] M. Jonker, R. Hofstede, A. Sperotto, and A. Pras, “Unveiling Flat Traffic on the Internet: An SSH Attack Case Study,” in *Proc. 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM’15)*, 2015, pp. 270–278.
- [55] M. Jonker and A. Sperotto, “Mitigating DDoS Attacks using OpenFlow-based Software Defined Networking,” in *Proc. of the 9th IFIP WG 6.6 International Conference on Autonomous Infrastructure, Management, and Security (AIMS’15)*, 2015, pp. 129–133.
- [56] M. Jonker and A. Sperotto, “Measuring Exposure in DDoS Protection Services,” in *Proc. of the 13th International Conference on Network and Service Management (CNSM’17)*, 2017, pp. 1–9.
- [57] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras, “Measuring the Adoption of DDoS Protection Services,” in *Proceedings of the 2016 ACM Internet Measurement Conference (IMC’16)*, 2016, pp. 279–285.
- [58] T. Jordan and P. Taylor, *Hacktivism and cyberwars: Rebels with a cause?* Routledge, 2004.
- [59] M. Karami and D. McCoy, “Understanding the Emerging Threat of DDoS-as-a-Service,” in *Proceedings of the 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats, LEET’13*, 2013.
- [60] S. Khandelwal, “602 Gbps! This May Have Been the Largest DDoS Attack in History,” <http://thehackernews.com/2016/01/biggest-ddos-attack.html>, January 2016.
- [61] A. King, “Corsaro,” 2012, <http://www.caida.org/tools/measurement/corsaro/>.
- [62] A. King, “Corsaro RS DoS Plugin,” 2012, https://www.caida.org/tools/measurement/corsaro/docs/plugins.html#plugins_dos.
- [63] T. King, C. Dietzel, J. Snijders, G. Doering, and G. Hankins, “BLACKHOLE Community,” RFC 7999 (Internet Standard), Internet Engineering Task Force, October 2016. (link)
- [64] S. Kitterman, “Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1,” RFC 7208 (Internet Standard), Internet Engineering Task Force, April 2014. (link)
- [65] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow, “AmpPot: Monitoring and Defending Against Amplification DDoS Attacks,” in *International Workshop on Recent Advances in Intrusion Detection (RAID’15)*, 2015, pp. 615–636.
- [66] B. Krebs, “DDoS Attack on Bank Hid \$900,000 Cyberheist,” <https://krebsonsecurity.com/2013/02/ddos-attack-on-bank-hid-900000-cyberheist/>, February 2013.

- [67] J. Krupp, M. Karami, C. Rossow, D. McCoy, and M. Backes, “Linking Amplification DDoS Attacks to Botnet Services,” in *Proc. of the 20th International Symposium on Research in Attacks, Intrusions, and Defenses (RAID '17)*, 2017, pp. 427–449.
- [68] W. Kumari and D. McPherson, “Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF),” RFC 5635 (Internet Standard), Internet Engineering Task Force, August 2009. (link)
- [69] E. L. and A. K., “Towards estimating the untapped potential: a global malicious DDoS mean capacity estimate,” *Journal of Cyber Policy*, vol. 2, no. 2, pp. 195–208, 2017.
- [70] J. Liu, “Firewall Rule: Global Traffic Management - SiteShield ACL,” <https://community.akamai.com/community/cloud-security/blog/2017/04/21/firewall-rule-global-traffic-management-siteshield-acl>, 2017, accessed: 2019-06-01.
- [71] G. Loukas and G. Öke, “Protection Against Denial of Service Attacks: A Survey,” *The Computer Journal*, vol. 53, no. 7, pp. 1020–1037, 2010.
- [72] P. Lutscher, N. Weidmann, M. E. Roberts, M. Jonker, A. King, and A. Dainotti, “At Home and Abroad: The Use of Denial-of-Service Attacks during Elections in Non-democratic Regimes,” *Journal of Conflict Resolution (JCR)*, July 2019.
- [73] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan, “Analyzing Large DDoS Attacks Using Multiple Data Sources,” in *Proceedings of the 2006 SIGCOMM Workshop on Large-scale Attack Defense (LSAD'06)*, 2006, pp. 161–168.
- [74] J. Markoff and N. Pelroth, “Firm Is Accused of Sending Spam, and Fight Jams Internet,” http://www.nytimes.com/2013/03/27/technology/internet/online-dispute-becomes-internet-snarling-attack.html?_r=0, accessed: 2019-06-01.
- [75] D. McDonald, “The Pentester’s Guide to Akamai,” 2013.
- [76] M. McKeay *et al.*, “The Q4 2016 State of the Internet / Security Report,” Akamai, 2017.
- [77] M. McKeay *et al.*, “The Q4 2017 State of the Internet / Security Report,” Akamai, 2017.
- [78] J. Mirkovic, S. Dietrich, D. Dittrich, and P. Reiher, *Internet Denial of Service: Attack and Defense Mechanisms (Radia Perlman Computer Networking and Security)*, 2004.

- [79] J. Mirkovic and P. Reiher, “A Taxonomy of DDoS Attack and DDoS Defense Mechanisms,” *SIGCOMM Computer Communications Review (CCR)*, vol. 34, no. 2, 2004.
- [80] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, “Inferring Internet Denial-of-service Activity,” *ACM Transactions on Computer Systems*, vol. 24, no. 2, pp. 115–139, 2006.
- [81] C. Morales, “NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us,” Mar 2018, <https://asert.arbornetworks.com/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/>.
- [82] G. C. Moura, R. d. O. Schmidt, J. Heidemann, W. B. de Vries, M. Muller, L. Wei, and C. Hesselman, “Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event,” in *Proceedings of the 2016 ACM Internet Measurement Conference (IMC’16)*, 2016.
- [83] A. Nixon and C. Camejo, “DDoS Protection Bypass Techniques,” Black Hat USA, 2013.
- [84] A. Noroozian, M. Korczyński, C. H. Gañan, D. Makita, K. Yoshioka, and M. van Eeten, “Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service,” in *Proc. of the 19th International Symposium on Research in Attacks, Intrusions, and Defenses (RAID ’16)*, 2016.
- [85] E. Nygren, R. K. Sitaraman, and J. Sun, “The Akamai Network: A Platform for High-performance Internet Applications,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, 2010.
- [86] P. Olson, “The Largest Cyber Attack In History Has Been Hitting Hong Kong Sites,” <https://www.forbes.com/sites/parmyolson/2014/11/20/the-largest-cyber-attack-in-history-has-been-hitting-hong-kong-sites/>, November 2014.
- [87] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, “BGPStream: A Software Framework for Live and Historical BGP Data Analysis,” in *Proc. of the 2016 Internet Measurement Conference (IMC ’16)*, 2016, pp. 429–444.
- [88] B. Overvest, M. Non, F. Ruesink, and R. Windog, “Risicorapportage Cyberveiligheid Economie 2018,” CPB Netherlands Bureau for Economic Policy Analysis, 2018.
- [89] G. Pack, J. Yoon, E. Collins, and C. Estan, “On filtering of DDoS attacks based on source address prefixes,” pp. 1–12, Aug 2006.
- [90] P. Paganini, “The hosting provider OVH continues to face massive DDoS attacks launched by a botnet composed at least of 150000 IoT devices.” <http://securityaffairs.co/wordpress/51726/cyber-crime/ovh-hit-botnet-iot.html>, September 2016.

- [91] R. A. Paulson and J. E. Weber, "Cyberextortion: an overview of distributed denial of service attacks against online gaming companies," *Issues in Information Systems*, vol. 7, no. 2, pp. 52–56, 2006.
- [92] J. Pescatore, "DDoS Attacks Advancing and Enduring: A SANS Survey," SANS, 2014.
- [93] M. Prince, "Technical Details Behind a 400Gbps NTP Amplification DDoS Attack," <https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/>, February 2014.
- [94] rdnl, "PAN, BBMRI-Omics and OpenINTEL are the winners of the Dutch Data Prize 2018," [http://researchdata.nl/en/news-and-agenda/news/news-item/?tx_news_pi1\[news\]=179](http://researchdata.nl/en/news-and-agenda/news/news-item/?tx_news_pi1[news]=179), 2018.
- [95] P. Richter, G. Smaragdakis, D. Plonka, and A. Berger, "Beyond Counting: New Perspectives on the Active IPv4 Address Space," in *Proceedings of the 2016 ACM Internet Measurement Conference (IMC'16)*, 2016.
- [96] S. Rose, R. Chandramouli, and A. Nakassis, "Information leakage through the domain name system," in *Proceedings of the Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH'09)*, 2009, pp. 16–21.
- [97] A. P. Rosiello, E. Kirda, C. Kruegel, and F. Ferrandi, "A Layout-Similarity-Based Approach for Detecting Phishing Pages," in *Proceedings of the 3rd International Conference on Security and Privacy in Communications Networks and the Workshops (SecureComm'07)*, 2007, pp. 454–463.
- [98] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *NDSS*, 2014.
- [99] J. J. Santanna, R. van Rijswijk-Deij, A. Sperotto, R. Hofstede, M. Wierbosch, L. Z. Granville, and A. Pras, "Booters - An Analysis of DDoS-as-a-Service Attacks," in *Proceedings of the 14th IFIP/IEEE International Symposium on Integrated Network Management (IM'15)*, 2015.
- [100] M. Sargent, J. Kristoff, V. Paxson, and M. Allman, "On the Potential Abuse of IGMP," *ACM SIGCOMM Computer Communication Review*, vol. 47, no. 1, 2017.
- [101] C. Shannon, D. Moore, K. Keys, M. Fomenkov, B. Huffaker, and k. claffy, "The Internet Measurement Data Catalog," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 35, no. 5, pp. 97–100, Oct 2005.
- [102] N. Sullivan, "DDoS Prevention: Protecting The Origin," <https://blog.cloudflare.com/ddos-prevention-protecting-the-origin/>, 2013, accessed: 2017-05-24.

- [103] D. Thomas, R. Clayton, and A. Beresford, “1000 days of UDP amplification DDoS attacks,” in *APWG Symposium on Electronic Crime Research (eCrime 2017)*, 2017.
- [104] O. van der Toorn, R. Hofstede, M. Jonker, and A. Sperotto, “A first look at HTTP (S) intrusion detection using NetFlow/IPFIX,” in *Proc. 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM’15)*, pp. 862–865.
- [105] R. van Rijswijk-Deij, M. Jonker, and A. Sperotto, “On the adoption of the elliptic curve digital signature algorithm (ECDSA) in DNSSEC,” in *Proc. of the 12th International Conference on Network and Service Management (CNSM’16)*, 2016, pp. 258–262.
- [106] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, “The internet of names: A dns big dataset,” *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, pp. 91–92, 2015.
- [107] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, “A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements,” *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 34, no. 6, pp. 1877–1888, 2016.
- [108] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, “DNSSEC and its potential for DDoS attacks: a comprehensive measurement study,” in *Proceedings of the 2014 ACM Internet Measurement Conference (IMC’14)*, 2014, pp. 449–460.
- [109] S. van Voorst, “Langdurig ddos-onderzoek wijst op 30.000 aanvallen per dag,” <https://tweakers.net/nieuws/131497/langdurig-ddos-onderzoek-wijst-op-30000-aanvallen-per-dag.html>, accessed: 2019-07-01.
- [110] R. Vasudevan, Z. Mao, O. Spatscheck, and J. van der Merwe, “Reval: A Tool for Real-time Evaluation of DDoS Mitigation Strategies,” in *In USENIX Annual Technical Conference*, Jun 2006, pp. 157–170.
- [111] T. Vissers, T. Van Goethem, W. Joosen, and N. Nikiforakis, “Maneuvering Around Clouds: Bypassing Cloud-based Security Providers,” in *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security (CCS’15)*, 2015, pp. 1530–1541.
- [112] A. Wang, A. Mohaisen, W. Chang, and S. Chen, “Delving into Internet DDoS Attacks by Botnets: Characterization and Analysis,” in *Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN’15)*, 2015, pp. 379–390.
- [113] S. Weagle, “Financial Impact of Mirai DDoS Attack on Dyn Revealed in New Data,” <https://www.corero.com/blog/797-financial-impact-of-mirai-ddos-attack-on-dyn-revealed-in-new-data.html>, February 2017.

- [114] F. Weimer, “Passive DNS replication,” in *FIRST Conference on Computer Security Incident*, 2005.
- [115] A. Welzel, C. Rossow, and H. Bos, “Delving into Internet DDoS Attacks by Botnets: Characterization and Analysis,” in *Proceedings of the 7th European Workshop on System Security (EuroSec’14)*, 2014, pp. 3:1–3:6.
- [116] M. Wullink, G. C. Moura, M. Müller, and C. Hesselman, “ENTRADA: A high-performance network traffic data streaming warehouse,” in *2016 IEEE/IFIP Network Operations and Management Symposium (NOMS’16)*, 2016, pp. 913–918.
- [117] S. Zander, L. L. Andrew, and G. Armitage, “Capturing Ghosts: Predicting the Used IPv4 Space by Inferring Unobserved Addresses,” in *Proceedings of the 2014 ACM Conference on Internet Measurement Conference (IMC’14)*, 2014.
- [118] S. T. Zargar, J. Joshi, and D. Tipper, “A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks,” *IEEE Communications Surveys Tutorials*, vol. 15, no. 4, 2013.
- [119] K. Zhang and D. Shasha, “Simple Fast Algorithms for the Editing Distance between Trees and Related Problems,” *SIAM Journal on Computing*, vol. 18, no. 6, pp. 1245–1262, 1989.
- [120] H. Zimmermann, “OSI reference model—The ISO model of architecture for open systems interconnection,” *IEEE Transactions on Communications*, vol. 28, no. 4, pp. 425–432, 1980.

About the Author

Mattijs was born in Alkmaar, The Netherlands. His passion for computer related matters developed at a young age and his first programming efforts trace back to *Atari BASIC* and a program cassette recorder. He earned a M.Sc. in Cyber Security from the University of Twente, Enschede, The Netherlands. In his master's thesis and the conference paper published alongside, Mattijs improved state-of-the-art techniques to detect brute-force attacks on security protocols such as SSH.

From 2014 to 2019, Mattijs pursued a Ph.D. degree in the *Design and Analysis of Communication Systems (DACS)* group at the University of Twente. During this time he worked on the D3 project, a national cyber security project funded by the Netherlands Organization for Scientific Research (NWO). The D3 project focused on protecting schools and other public organizations against Distributed Denial-of-Service attacks. He was also involved in the European Union FP7 SALUS project, the goal of which was to design, implement and evaluate a next generation communication network concept for Public Protection and Disaster Relief (PPDR) agencies. Shortly before finishing his Ph.D., Mattijs started working on the EU Horizon 2020 CONCORDIA project, the goal of which is to establish an user-centric EU-integrated cyber security ecosystem for digital sovereignty in Europe

Besides his work, Mattijs enjoys traveling, photography and clearing his mind at the gym. He is also a beginning scuba diver.

More information about Mattijs can be found on his Web site at <https://mattijsjonker.com>. He can also be reached via his private e-mail address: me@mattijsjonker.com.

List of Publications

Below is a list of peer-reviewed academic publications that Mattijs has authored or co-authored. They are listed in reverse chronological order.

- P. Lutscher, Nils Weidmann, Molly E. Roberts, M. Jonker, A. King and A. Dainotti. *At Home and Abroad: The Use of Denial-of-Service Attacks during Elections in Non-democratic Regimes*. Journal of Conflict Resolution (JCR; July 2019) [72].
- M. Jonker, A. Pras, A. Dainotti and A. Sperotto. *A First Joint Look at DoS Attacks and BGP Blackholing in the Wild*. In proceedings of the 2018 ACM Internet Measurement Conference (IMC'18). Boston, Massachusetts, USA [52].
- N. Hinze, M. Nawrocki, M. Jonker, A. Dainotti, T.C. Schmidt, and M. Wählich. *On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP*. In proceedings of ACM SIGCOMM 2018 on Posters and Demos. Budapest, Hungary [45].
- M. Jonker and A. Sperotto. *Measuring Exposure in DDoS Protection Services*. In proceedings of the 13th International Conference on Network and Service Management (CNSM'17). Tokyo, Japan [56].
- M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto and A. Dainotti. *Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem*. In proceedings of the 2017 ACM Internet Measurement Conference (IMC'17). London, United Kingdom [51].
- R. Hofstede, M. Jonker, A. Sperotto and A. Pras. *Flow-Based Web Application Brute-Force Attack and Compromise Detection*. Journal of Network and Systems Management (JONS). Volume 1, Issue 4 (August 2017) [46].
- M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre and A. Pras. *Measuring the Adoption of DDoS Protection Services*. In proceedings of the 2016 ACM Internet Measurement Conference (IMC'16). Santa Monica, California, USA [57].
- R. van Rijswijk-Deij, M. Jonker and A. Sperotto. *On the Adoption of the Elliptic Curve Digital Signature Algorithm (ECDSA) in DNSSEC*. In proceedings of the 12th International Conference on Network and Service Management (CNSM'16). Montréal, Québec, Canada [105].
- R. van Rijswijk-Deij, M. Jonker, A. Sperotto and A. Pras. *A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements*. IEEE Journal on Selected Areas in Communication (JSAC), Volume 34, Issue 6 (June 2016) [107].
- R. van Rijswijk-Deij, M. Jonker, A. Sperotto and A. Pras. *The Internet of Names: A DNS Big Dataset – Actively Measuring 50% of the Entire DNS Name Space, Every Day*. SIGCOMM Computer Communications Review (CCR). Volume 45, Issue 4 (October 2015) [106].

- M. Jonker and A. Sperotto. *Mitigating DDoS Attacks using OpenFlow-based Software Defined Networking*. In Proceedings of the 9th IFIP WG 6.6 International Conference on Autonomous Infrastructure, Management, and Security (AIMS'15). Ghent, Belgium [55].
- M. Jonker, R. Hofstede, A. Sperotto and A. Pras. *Unveiling Flat Traffic on the Internet: An SSH Attack Case Study*. In proceedings of the 14th IFIP/IEEE International Symposium on Integrated Network Management (IM'15). Ottawa, Ontario, Canada [54].
- O. van der Toorn, R. Hofstede, M. Jonker and A. Sperotto. *A First Look at HTTP(S) Intrusion Detection using NetFlow/IPFIX*. In proceedings of the 14th IFIP/IEEE International Symposium on Integrated Network Management (IM'15). Ottawa, Ontario, Canada [104].

ISBN: 978-90-365-4868-7

ISSN: 2589-7721

DSI Ph.D. Thesis Series No. 19-018

DOI: 10.3990/1.9789036548687