

A First Joint Look at DoS Attacks and BGP Blackholing in the Wild

Mattijs Jonker
University of Twente
m.jonker@utwente.nl

Aiko Pras
University of Twente
a.pras@utwente.nl

Alberto Dainotti
CAIDA / UC San Diego
alberto@caida.org

Anna Sperotto
University of Twente
a.sperotto@utwente.nl

ABSTRACT

BGP blackholing is an operational countermeasure that builds upon the capabilities of BGP to achieve DoS mitigation. Although empirical evidence of blackholing activities are documented in literature, a clear understanding of how blackholing is used in practice when attacks occur is still missing.

This paper presents a first joint look at DoS attacks and BGP blackholing in the wild. We do this on the basis of two complementary data sets of DoS attacks, inferred from a large network telescope and DoS honeypots, and on a data set of blackholing events. All data sets span a period of three years, thus providing a longitudinal overview of operational deployment of blackholing during DoS attacks.

CCS CONCEPTS

• **Networks** → Denial-of-service attacks; Network measurement; Network management; Routing protocols; • **Security and privacy** → Security services;

KEYWORDS

Denial-of-Service; DDoS Mitigation; BGP; Blackholing

ACM Reference Format:

Mattijs Jonker, Aiko Pras, Alberto Dainotti, and Anna Sperotto. 2018. A First Joint Look at DoS Attacks and BGP Blackholing in the Wild. In *2018 Internet Measurement Conference (IMC '18)*, October 31–November 2, 2018, Boston, MA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3278532.3278571>

1 INTRODUCTION

Volumetric Denial-of-Service (DoS) attacks have rapidly increased in frequency and intensity over the last years. In previous work, we found an average of thirty thousand attacks daily, with intensities ranging from a mere nuisance to severe [1]. Thanks to so-called Booters [2], DoS has also become available “as-a-Service”, allowing the layman to launch attacks powerful enough to saturate *1-10 Gbps* links. The full potential of attacks has arguably yet to be seen and Leverett et al. [3] estimate the upper bound of distributed reflection and amplification attacks to be above *100 Tbps*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

IMC '18, October 31–November 2, 2018, Boston, MA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5619-0/18/10...\$15.00

<https://doi.org/10.1145/3278532.3278571>

The fight against DoS attacks has prompted the development of diverse mitigation techniques. Examples are cloud-based DDoS Protection Services [4], which use traffic diversion to third-party data centers that “cleanse” traffic; on-site, in-line appliances (e.g., those offered by Netscout Arbor [5] and Radware [6]); BGP Flowspec [7] or BGP blackholing.

This paper focuses on BGP blackholing, an operational countermeasure that builds upon the capabilities of the Border Gateway Protocol (BGP) to achieve DoS mitigation. BGP blackholing is implemented using the BGP *communities attribute* [8], a BGP extension that enables passing additional information to BGP peers [9]. BGP blackholing makes use of a specific set of BGP community tags to request an upstream provider (ISP) or IXP to filter, i.e., null-route traffic to a specific destination prefix (the one of the victim) [10].

Although empirical evidence of blackholing activities is documented in literature [11], a clear understanding about how BGP blackholing is used in practice when attacks occur is still missing. The goal of this paper is to provide a first joint look at DoS attacks and BGP blackholing in the wild. To this end, we rely on two data sets of DoS attacks and one of blackholing events, all spanning a little over three years (1100 days). To the best of our knowledge, this is the first large-scale empirical observation of DoS events and corresponding blackholing mitigation. Our main findings are:

- Mitigation via blackholing happens within minutes. Our analysis shows that 44% of the attacks for which blackholing is put in place are mitigated within one minute, and 85% within ten minutes.
- A significant fraction of blackholing events show blackholing is still in place hours after the end of the attack, which raises the question if the remedy is in some cases worse than the disease, as any service and system in the blackholed prefix might experience lack of connectivity or it needs to rely on alternative routes for longer than necessary.
- 13% of the blackholing events in our data set is related to attacks with very low intensity, specifically *3 Mbps* or less. This finding has two main implications. First, it indirectly confirms the findings of the seminal paper of Moore et al. [12], by explicitly linking low-intensity backscatter to actual DoS mitigation. The second implication is operational. BGP Blackholing is a coarse-grained mitigation strategy. One could imagine that blackholing is therefore only used for large attacks as a last resource, that is, if other fine-grained solutions (e.g., scrubbing, flowspec) do not work. Our analysis shows that this is not the case, raising the question of what is the minimal effort needed by an attacker to trigger such a drastic countermeasure.

The remaining of this paper is organized as follows. Sections 2 and 3 present the data sets used in our analysis, and our results, respectively. In Section 4 we discuss related work. Finally, in Section 5 we briefly discuss limitations and future work.

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 1: Denial-of-Service data from UCSD-NT and AmpPot for March 1, 2015 – March 5, 2018. We find 28.14 M attacks, targeting 8.58 M unique IP addresses.

2 DATA SETS

In this paper, we consider two DoS attack events data sets and one data set of BGP blackholing events. All data sets cover the same period, from March 1, 2015 through March 5, 2018.

2.1 DoS Attack Events

The DoS data sets contain various attack types, measured by established and complementary data sources.

Randomly and Uniformly Spoofed Attacks – The first data set on DoS attacks is inferred from backscatter packets that reach the UCSD Network Telescope [13] (UCSD-NT). The UCSD-NT is a largely-unused but routed /8 network operated by University of California, San Diego. It passively collects unsolicited traffic resulting from, among others, scans, misconfigurations, and backscatter from Denial-of-Service attacks. The UCSD-NT covers approximately $1/256$ of the IPv4 address space. This means that a randomly and uniformly selected IPv4 address has an approximate probability of $1/256$ to fall within UCSD-NT’s address space. Randomly and uniformly spoofed attacks are often visible at the UCSD-NT as these attacks typically involve backscatter to a substantial number [1, 12] of spoofed IPv4 addresses. To infer attacks we use the classification methodology described by Moore et al. [12]. For each attack we register, among others: the attack’s target, *i.e.*, intended victim – apparent from the backscatter packets; the attack’s (observed) beginning and end times; and a measure of attack intensity based on backscatter packet rate.¹ Further details on the implementation can be found in Jonker et al. [1].

Reflection and Amplification Attacks – The second data set on DoS attacks is inferred in honeypots running AmpPot [14]. In reflection and amplification attacks, requests with a specifically spoofed source address are used to trigger reflectors to send unrequested response packets. The address is set to be that of the intended victim and the responses are typically considerably larger than the requests (*i.e.*, there is amplification). AmpPot emulates various protocols known to be abused in this type of attack, such as NTP, DNS and CharGen [15]. During an attack, the attacker sends requests – apparently coming from the intended victim – to AmpPot. AmpPot records these requests and registers various information about each attack, such as: the target – apparent from the source address spoofed in the requests; a measure of attack intensity based on the request rate; and the attack’s (observed) beginning and end times. We use data from 24 AmpPot instances.² It has been shown that this number of AmpPot instances is sufficient to register most reflection

¹As not all attack traffic leads to backscatter, this intensity forms a lower bound.

²The US houses 11 instances, 8 are in Europe, 4 are in Asia and 1 is in Australia.

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

Table 2: Blackholing data set inferred from public BGP data for March 1, 2015 – March 5, 2018. We infer 1.3 M blackholing events, involving 146193 prefixes.

attacks on the Internet. Further details on AmpPot can be found in Krämer et al. [14].

The UCSD-NT data set includes spoofing attacks that directly target the victim. The AmpPot data set, differently, reports on indirect (reflected) attacks. As such, the two data sets complement each other. The data sets, however, do not cover attacks in which packets are sent without any form of source IP address spoofing. Table 1 summarizes the data sets in terms of attack events, targets and involved ASNs.

2.2 Blackholing Events

We obtain a data set of inferred blackholing events from publicly available BGP routing data, using a measurement system that we implemented on the basis of the methodology described by Giotsas et al. [11].

Public BGP data – We use data from two projects that offer public BGP data: (1) University of Oregon’s *RouteViews Project (RV)* [16]; and (2) RIPE NCC’s *Routing Information Service (RIS)* [17]. Both these projects gather Internet routing data from globally dispersed collectors that peer with one or multiple routers.³

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. [11] created a dictionary of such communities by applying natural language processing to resources where blackholing communities are likely to be documented (*e.g.*, in Internet Routing Registry (IRR) records). 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).⁴

Inferring Blackholing Events – We implemented a measurement system in Python that utilizes *pyBGPStream*, a Python interface to the *BGPStream* framework for BGP data analysis [18]. Because of our focus, we do not consider prefixes less specific than a /24, since these are not commonly blackholed [11, 19]. We do infer blackholing activity incrementally, by analyzing BGP updates, and do not parse a full RIB dump at the beginning of the observation period.⁵ To create our data set, we analyze data from 36 BGP collectors.⁶ Each event in the data set contains, most notably: the blackholed prefix, a start time (*i.e.*, activation time), an optional end time (*i.e.*, deactivation time), a list of collectors on which prefix-related activity was observed⁷, and the matched communities.

³*Packet Clearing House* also provides public BGP data; we do not use these, primarily due to lack of support in the *BGPStream* framework.

⁴The dictionary contains the majority of BGP blackholing communities, but it is not necessarily complete due to methodological limitations [11].

⁵Consequently, we will miss blackholing events that started before March 2015.

⁶Not all blackholing announcements propagate as far as public BGP collectors, meaning that we cannot possibly infer all blackholing events [11].

⁷Blackholing activity is considered related if it (partially) overlaps in time. An event’s activation and deactivation are set to the minimum and maximum BGP record

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 3: 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.

Table 2 summarizes our data set. 34 of the 36 collectors we consider see at least one blackholing event in the measurement period.⁸ 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.⁹

3 BLACKHOLED ATTACKS

We 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 a blackholing event’s prefix, and the attack’s start time to precede the blackholing event’s activation in time (of at most 24 hours).¹⁰

Table 3 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.**

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 look 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) its 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.

timestamps encountered in BGP announcements and withdrawals. 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, or through a prefix withdrawal. We presume consistent propagation characteristics between announcements and withdrawals.

⁸The 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*).

⁹This can occur if the event is inferred from BGP events on multiple collectors.

¹⁰We will show that blackholing is often triggered well within the hour following an attack’s start time.

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 4: 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%).

Interestingly, for the 447.6 k attacks jointly launched against the same target (Table 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 [1]. Table 4 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 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.**




Figure 1: 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 1 shows the time it takes for blackholing to be activated. For any blackholed attack in the data sets, we analyze the delay between the start of the attack and the start of the associated blackholing event.¹¹ For joint

¹¹BGP 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.

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.^{12,13} 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 2: 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 2 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.¹⁴ 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 sort of self-inflicted DoS. Later in this section we provide some results about the potential impact of blackholing on different type of infrastructure.

We see evidence that less intense attacks are also mitigated. 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 3 shows the overall distribution of intensities in the UCSD-NT data set, as well as for blackholed attacks only. 64.6% of blackholed attacks (gray curve) have




Figure 3: 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).

an intensity not greater than $100 pps_{max}$, which corresponds to an approximate attack traffic volume of 300 Mbps.¹⁵ 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.¹⁶ More importantly, a non-negligible percentage of blackholed attacks have 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. [12]) are relevant, since they trigger mitigation. Finally, this result underpins the validity of the surprisingly large number of DoS attacks we discovered in a recent work [1], contributing to the bigger picture, and it provides a reference threshold to be used in the context of monitoring and situational awareness.

The analysis of blackholed reflection attacks yields similar results. The AmpPot data set contains an intensity measure (rps_{avg}), expressed in terms of the average number of requests per second, *e.g.*, DNS queries.¹⁷ 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 4 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

¹²In 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 (Section 2.1). However, based on our observations of randomly spoofed attacks and reflection attacks, joint attacks are relatively rare.

¹³We analyzed the start time differences between attack components of the 18.4k joint blackholed attacks in our data (Table 3) and find that 85.54% see the attack start spaced less than 40 minutes apart.

¹⁴Blackholing “truncates” the attack end times in UCSD-NT data, which is why we do not analyze deactivation delays for randomly spoofed attacks.

¹⁵We assume 1500-byte packets and account for UCSD-NT’s $1/256$ address space coverage, *i.e.*, observing 1 backscatter packet for every 256 uniformly spoofed packets.

¹⁶In previous work we showed that stronger attacks lead to quicker outsourcing to DDoS Protection Services – another form of mitigation [1].

¹⁷AmpPot honeypots are part of a larger set of amplifiers. The attack intensity depends on all amplifiers involved, and honeypots cannot know the extent of involvement. By a best-effort guess, the number of amplifiers will not vary significantly among attacks for the same reflection protocol [1]. We thus consider the intensity per protocol.




Figure 4: 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.

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 18.4k joint blackholed attacks in our data (Table 3). 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.

The blackholing communities we observe reflect actual traffic filtering. Figure 5 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-directed traffic, including the packets that trigger backscatter. Consequentially, the attack end time observed through backscatter may not reflect the

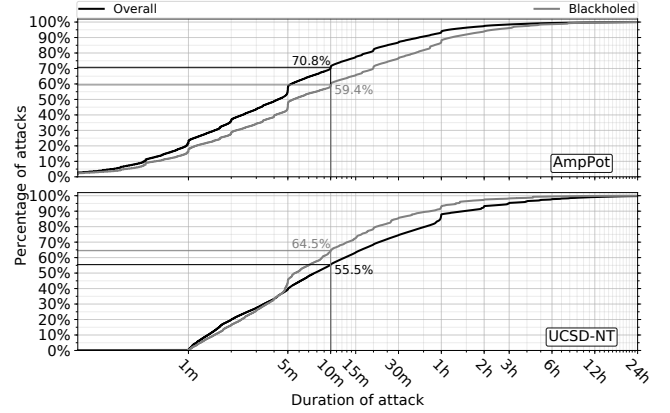


Figure 5: the attack duration distributions for all attacks (black curves) and blackholed attacks (gray curves) in the amppot data (upper plot) and the ucspd-nt data (lower plot). We find that for randomly spoofed attacks, the average duration drops, which, given the attack-inference methodology, is indicative that blackholing is effectively stopping (at least part) of victim-directed traffic.

start	#days	type	#names	#IPs
2015-03-01	1100	Web	228.1 M	33.5 M
2017-01-22	407	Mail (MX)	38.76 M	4.73 M
		DNS (NS)	7.62 M	1.54 M

Table 5: Active DNS measurement data for Web sites, mail exchangers and name servers. We observe a total of 228.1M Web sites for March 1, 2015 – March 5, 2018 (1100 days), and 38.76 M and 7.62 M unique mail exchanger and name server names for January 22, 2017 – March 5, 2018 (407 days).

type	#names		ratio (%)
	all	no-alt	
Web	754073 (0.33%)	658704	87.4
Mail (MX)	154200 (0.40%)	151117	98.0
DNS (NS)	9994 (0.13%)	9858	98.6

Table 6: Web sites, mail and name servers hosted in blackholed prefixes. For the relatively small percentages of associations that we find, 87.4 to 98.6% do not have an alternative, non-blackholed IP address.

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.

Loss of service may affect Web sites, mail and name server infrastructure. 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 explore the impact blackholing may have on the availability of services by considering data from OpenINTEL¹⁸. OpenINTEL is an active DNS measurements platform [20] that measures daily snapshots of the DNS by querying all domain names under Top-Level Domains (TLDs) for their Resource Records (RRs). This includes IP addresses of: (i) www labels, (ii) mail exchanger (MX), and (iii) authoritative name servers (NS). We use these records to map Web sites, mail and name servers to blackholing events.¹⁹

We use data for the three generic TLDs: .com, .net, and .org, which cover 50% of the global namespace [21]. Table 5 summarizes the data. We note that the Web site data spans the full DoS and blackholing data sets, but MX and NS data is shorter as the functionality to resolve these records was added to OpenINTEL later (on January 22, 2017).

Table 6 summarizes the blackholing correlation. 754k Web sites map to blackholing (0.33% of 228.1M). Of unique MX and NS names, 154k (0.40% of 38.76M) and 9994 (0.13% of 7.62M) map to blackholed prefixes.²⁰

Infrastructure can be redundantly hosted, *i.e.*, have multiple IP addresses. We investigate this by studying the presence of non-blackholed IP address records and find that, respectively, 87.4%, 98.0% and 98.6% of the names found (*cf.*, *ratio* in Table 6) do not have an alternative IP address at the time of blackholing. It follows that these services may become – and remain for extended time – unavailable when blackholing is left active and no IP address change takes place. We note that Mail Transfer Agents (MTAs) typically try to resend mails for days, meaning that an unreachable mail exchanger may incur a (hefty) delay rather than a full loss of service. When name servers cannot be reached, the domain names for which they are authoritative may become unresolvable – something also influenced by DNS redundancy and caching mechanisms (*e.g.*, TTL settings). A thorough characterization of loss of service requires extensive additional measurement – some to be performed on-the-fly while the event is happening, which we started instrumenting and leave as future work.

4 RELATED WORK

Several approaches to DoS mitigation have been proposed in literature [22, 23]. However, it is only with the analysis in Jonker et al. [4] that a first characterization of the adoption of DoS mitigation solutions was measured at scale. Giotsas et al. [11] present a comprehensive characterization of BGP blackholing activity, based on BGP data. Dietzel et al. [24] and Chatzis et al. [25] emphasize that IXPs play a key role in deploying blackholing. Our contribution focuses instead on correlating DoS attacks and blackholing events.

DoS attacks have been the subject of detailed studies [26]. A large-scale, measurement-based characterization of DoS attacks at the macroscopic scale is only given, however, in Jonker et al. [1], with an analysis of reflection and randomly spoofed attacks over a two-year period. The paper also retraces the major steps in DoS analysis, from the seminal paper by Moore et al. [12]. DoS characterization

has been carried on based on diverse data sources. While the work in Moore et al. focuses on the analysis of backscatter, Krämer et al. [14] and Thomas et al. [27] focus on DoS attacks as seen from a set of amplification honeypots. Santanna et al. [2] and Krupp et al. [28] analyze DoS attacks generated by Booters, while Wang et al. [29] analyze Botnet-related attacks. The focus of this paper is not on DoS attacks per se, but on the relation between DoS attacks and blackholing.

5 CONCLUSIONS

This study compares, on a global scale, DoS attacks with BGP blackholing events, revealing 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. Our preliminary results also highlight that, to further understand the impact of blackholing, more data is needed, *e.g.*, on-the-fly DNS measurements triggered once a blackhole is announced. Such measurements will shed light on aspects such as which networks are fully taken offline by a blackhole and which services benefit from blackholing but subsequently migrate to a different IP address to ensure service continuity. Finally, this study contributes to a better understanding of the whole DoS ecosystem: (i) it validates and reinforces some findings from our previous work [1], and (ii) it adds other pieces to the puzzle of the bigger picture of DoS attacks (attacks, defenses, impact, etc.).

ACKNOWLEDGMENTS

This work is part of the NWO: D3 project, which is funded by the Netherlands Organization for Scientific Research (628.001.018). This research was made possible by OpenINTEL, a joint project of the University of Twente, SURFnet, and SIDN. This material is based on research sponsored by Air Force Research Laboratory under agreement number FA8750-18-2-0049. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575. This work was supported by National Science Foundation grant CNS-1730661. We are grateful to Christian Rossow and Johannes Krupp for sharing AmpPot data. We thank our shepherd and the anonymous reviewers for their valuable feedback.

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

¹⁸<https://openintel.nl/>

¹⁹The existence of an A RR for the www label is taken as a Web site indicator.

²⁰Multiple domains may share the same infrastructure.

REFERENCES

- [1] 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)*, pages 100–113, 2017.
- [2] 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 *Proc. of the 14th IFIP/IEEE International Symposium on Integrated Network Management (IM '15)*, 2015.
- [3] Eireann L. and Aaron K. Towards estimating the untapped potential: a global malicious DDoS mean capacity estimate. *Journal of Cyber Policy*, 2(2):195–208, 2017.
- [4] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. Measuring the Adoption of DDoS Protection Services. In *Proc. of the 2016 ACM Internet Measurement Conference (IMC '16)*, pages 279–285, 2016.
- [5] Arbor solutions. <https://www.netscout.com/arbor>.
- [6] DDoS Prevention Services: Multi Layered DDoS Security Solutions. <https://www.radware.com/solutions/security/>.
- [7] Implementing BGP Flowspec. https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-3/routing/configuration/guide/b_routing_cg53xasr9k/implementing_bgp_flowspec.pdf. accessed: 2018-05-01.
- [8] B. Donnet and O. Bonaventure. On BGP Communities. *SIGCOMM Computer Communications Review (CCR)*, 38(2), 2008.
- [9] R. Chandra, P. Traina, and T. Li. BGP Communities Attribute. RFC 1997 (Internet Standard), August 1996.
- [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: 2018-05-01.
- [11] 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)*, pages 1–14, 2017.
- [12] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage. Inferring Internet Denial-of-service Activity. *ACM Transactions on Computer Systems*, 24(2):115–139, 2006.
- [13] UCSD Network Telescope (UCSD-NT). http://www.caida.org/projects/network_telescope/.
- [14] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, k. Yoshioka, and C. Rossow. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proc. of the International Workshop on Recent Advances in Intrusion Detection (RAID '15)*, pages 615–636, 2015.
- [15] C. Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proc. of the 2014 Network and Distributed System Security Symposium (NDSS '14)*, 2014.
- [16] University of Oregon Route Views Project. <http://www.routeviews.org>.
- [17] RIPE NCC Routing Information Service (RIS). <http://https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>.
- [18] 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)*, pages 429–444, 2016.
- [19] W. Kumari and D. McPherson. Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF). RFC 5635 (Internet Standard), August 2009.
- [20] 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)*, 34(6):1877–1888, 2016.
- [21] The Verisign Domain Name Industry Brief. https://www.verisign.com/en_US/domain-names/dnib/index.xhtml, 2018. accessed: 2018-05-01.
- [22] G. Loukas and G. Öke. Protection Against Denial of Service Attacks: A Survey. *The Computer Journal*, 53(7):1020–1037, 2010.
- [23] 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*, 15(4), 2013.
- [24] C. Dietzel, A. Feldmann, and T. King. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *Proc. of the 17th International Conference on Passive and Active Measurement (PAM '16)*, pages 319–332, 2016.
- [25] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is More to IXPs Than Meets the Eye. *SIGCOMM Computer Communications Review (CCR)*, 43(5), 2013.
- [26] Jelena Mirkovic and Peter Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *SIGCOMM Computer Communications Review (CCR)*, 34(2), 2004.
- [27] D. Thomas, R. Clayton, and A. Beresford. 1000 days of UDP amplification DDoS attacks. In *APWG Symposium on Electronic Crime Research (eCrime 2017)*, 2017.
- [28] 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)*, pages 427–449, 2017.
- [29] A. Wang, A. Mohaisen, W. Chang, and S. Chen. Delving into Internet DDoS Attacks by Botnets: Characterization and Analysis. In *Proc. of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN '15)*, pages 379–390, 2015.