Down the Black Hole: Dismantling Operational Practices of BGP Blackholing at IXPs

Marcin Nawrocki
marcin.nawrocki@fu-berlin.de
Freie Universität Berlin
Germany

Jeremias Blendin
jeremias.blendin@de-cix.net
DE-CIX
Germany

Christoph Dietzel
christoph@mpi-inf.mpg.de
DE-CIX / MPI for Informatics
Germany

Thomas C. Schmidt
t.schmidt@haw-hamburg.de
HAW Hamburg
Germany

Matthias Wählisch
m.waehlisch@fu-berlin.de
Freie Universität Berlin
Germany

ABSTRACT

Large Distributed Denial-of-Service (DDoS) attacks pose a major threat not only to end systems but also to the Internet infrastructure as a whole. Remote Triggered Black Hole filtering (RTBH) has been established as a tool to mitigate inter-domain DDoS attacks by discarding unwanted traffic early in the network, e.g., at Internet eXchange Points (IXPs). As of today, little is known about the kind and effectiveness of its use, and about the need for more fine-grained filtering.

In this paper, we present the first in-depth statistical analysis of all RTBH events at a large European IXP by correlating measurements of the data and the control plane for a period of 104 days. We identify a surprising practice that significantly deviates from the expected mitigation use patterns. First, we show that only one third of all 34k visible RTBH events correlate with indicators of DDoS attacks. Second, we witness over 2000 blackhole events announced for prefixes not of servers but of clients situated in DSL networks. Third, we find that blackholing on average causes dropping of only 50% of the unwanted traffic and is hence a much less reliable tool for mitigating DDoS attacks than expected. Our analysis gives also rise to first estimates of the collateral damage caused by RTBH-based DDoS mitigation.

CCS CONCEPTS

• **Security and privacy** → **Denial-of-service attacks**; • **Networks** → *Public Internet*.

KEYWORDS

DDoS, BGP, RTBH, Collateral Damage

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

ACM ISBN 978-1-4503-6948-0/19/10…$15.00
https://doi.org/10.1145/3355369.3355593

ACM Reference Format:
Marcin Nawrocki, Jeremias Blendin, Christoph Dietzel, Thomas C. Schmidt, and Matthias Wählisch. 2019. Down the Black Hole: Dismantling Operational Practices of BGP Blackholing at IXPs. In *Internet Measurement Conference (IMC '19), October 21–23, 2019, Amsterdam, Netherlands.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3355369.3355593

1 INTRODUCTION

The Border Gateway Protocol (BGP) is used to exchange IP prefix reachability information between Autonomous Systems (ASes) to form the global Internet. Yet, one BGP application has the opposite effect in practice: Signaling Remotely Triggered Black Hole filtering (RTBH) through BGP requests a neighboring AS to discard traffic destined towards an owned IP prefix. The most prominent and well-established use case for RTBH filtering is the mitigation of volumetric Distributed Denial-of-Service (DDoS) attacks. Recent attacks peak beyond multiple Tbps (Terabit per second) [28]. DDoS attacks build upon simple to exploit IP address spoofing [8, 9] in combination with amplification characteristics of network protocols such as NTP, DNS, or cLDAP [3, 15]. These attacks deplete network bandwidth to suppress legitimate traffic towards a destination IP. In consequence, a network or web service is not reachable anymore. Still, DDoS attacks do not only cause damage at the attacked system itself, but can also overwhelm the infrastructure of intermediate or upstream networks [37]. Such collateral damage often impairs common customers badly.

Intermediate ASes mitigate the collateral damage of DDoS traffic passing through their infrastructure by signaling RTBHs to their neighbors that specifically cover the target address of the DDoS attack. Thereby, volumetric attack traffic is dropped before it reaches the final destination and alleviate the damage to the network infrastructure under attack. Internet exchange points (IXP) are particularly well suited for this kind of prevention, since they provide a convergence point where hundreds of ASes meet and exchange inter-domain traffic [1, 5].

RTBH filtering is a light-weight and easy to use tool. It is widely deployed and can be highly effective, that is why RTBH is a well established reactive DDoS mitigation technique today [10]. On the downside, RTBH is a coarse granular mechanism that drops all traffic to a specific prefix, and does not provide information about the attack traffic while it is ongoing. Therefore, advanced alternatives such as ACL filters [8], BGP FlowSpec [12, 24, 40],

and Advanced Blackholing [6] have been introduced. Yet, RTBH continues to play a significant role in DDoS mitigation.

Understanding RTBH's operational intricacies and use cases as well as its traffic patterns and efficacy are crucial for understanding the effectiveness and success of RTBH and for the evaluation of its alternatives. Furthermore, an in-depth investigation can help to uncover issues with the established ways of using RTBH on the Internet. However, understanding the operational practices as well as the corresponding traffic patterns at large scales is limited in both academia and industry. A number of publications on RTBH traffic patterns describe representative investigations of single events [5, 6], but no large-scale, statistical analysis with in-depth empirical evidence exists.

We start with exploring the operational practices of RTBH at a large IXP and separate RTBH by their inferred use case. Thereby, DDoS attack mitigation RTBHs can be separated from other use cases and investigated in detail. We analyze the traffic patterns of DDoS RTBHs and gain thorough insights how these are connected to operational practices. To our surprise, we find use patterns and deployment of RTBH in the wild that differ widely from common expectations.

Our contributions are as follows:

(1) A description and characterization of RTBH use cases based on the literature as well as industry expert interviews
(2) We collect and analyze unsampled RTBH data over a period of three months and classify RTBH events by their use cases
(3) We uncover the statistical efficacy of RTBH traffic dropping over a large data set of 590 million sampled flows
(4) We present a detailed correlation analysis between DDoS attacks on the data plane and RTBH signaling on the control plane
(5) We provide insights into the detrimental effects of dropping attack traffic completely and quantify the beneficial traffic

This paper is organized as follows. We present the understanding of RTBH use cases and literature in Section 2. Section 3 introduces our control and data plane data set. We investigate which features of RTBH and to what extent they are used in Section 4. Empirical evidence on the traffic characteristics of DDoS mitigation RTBHs is presented in Section 5 followed by an investigation on collateral damage of RTBH filtering in Section 6. We discuss our findings on the background of discussions with industry experts in Section 7 before drawing a conclusion in Section 8.

2 BACKGROUND AND USE CASES

RTBH filtering is thought to be originally conceived to mitigate DDoS attacks on the Internet. Its low operational overhead to signal blacklisting makes it attractive for other use cases as well. In this section, we introduce RTBH as a tool for protecting infrastructure from DDoS attacks. Furthermore, we identify the use of RTBH in the context of prefix squatting protection and content blocking as well. Finally, we describe the expected blackhole characteristics for every use case based.

2.1 RTBH Primer

Remotely Triggered Blackholing uses BGP to signal blackholes, in contrast to other blackholing approaches such as access control lists. To start (or stop) a blackhole at IXPs, a member sends a BGP announcement (or withdrawal) to the IXP route server. The route server distributes the blackhole route to all or a subset of IXP peers, including a specific next hop IP address (*i.e.,* the blackhole). It is worth noting that any peer applies local BGP policies on the received blackhole route to decide whether to accept or reject the received blackhole route, as the peer does for any other route. Based on this decision, subsequent data that matches the blackhole route will be forwarded to the IXP infrastructure and dropped (see Figure 1).

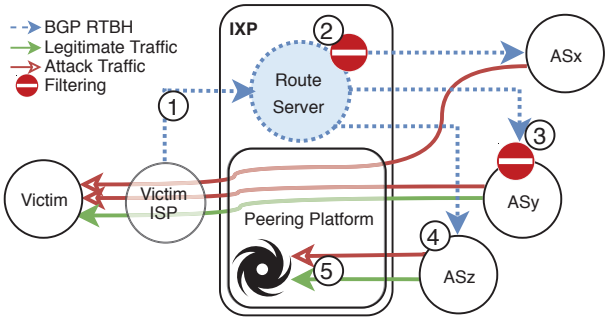

Figure 1: Remotely Triggered Blackholing (RTBH) at IXPs: ❶ RTBH announcement via BGP, ❷ Propagation filter, ❸ BGP policy rejects RTBH route, ❹ BGP policy accepts RTBH route, ❺ Packet drop.

A known drawback of RTBH is the collateral damage due to the rather coarse granularity of destination IP prefixes [12]. RTBH drops *all* traffic towards the prefix under protection, *i.e.,* legitimate traffic as well as attack traffic, because it cannot distinguish services on the transport layer. We will discuss this further in the next sections. Misusage of RTBH in combination with BGP signaling such as BGP blackjacks [26] is out of scope of this paper.

2.2 Infrastructure Protection

RTBH was designed to prevent forwarding of unwanted traffic [21, 46], *e.g.,* (*i*) attack traffic (DoS), (*ii*) incoming scan traffic [7], or (*iii*) Internet background radiation [34]. For the latter two, the traffic volume is comparatively small and operational best practices such as firewalls and static ACL filters [8] are adequate solutions. In contrast, today's terabit-level DDoS attacks are a serious threat to the operation of Web services [20, 28], and even challenge the Internet backbone infrastructure [37]. To alleviate the negative impact on the Internet infrastructure, RTBH is used as a cheap and convenient technique to filter unwanted traffic at intermediate network nodes [10]. Such a central location to blackhole unwanted traffic are IXPs [5]. Traffic of hundreds of ASes can be dropped or filtered on the IXP switching platform [6], making IXPs a good vantage point for this kind of studies.

For the usage of RTBH at an IXP, we expect a significant rise of inter-domain traffic volume, seen by a member. In reaction, this member will send most likely a `/32` RTBH. Note, RTBHs are announced and withdrawn constantly by the victims to gather attack status information—if the traffic is discarded no telemetry data is available [10]. Thus, our assumption is to observe a temporally

Table 1: Literature-based expected characteristics of RTBHs by use case.

Use Case	Trigger	Prefix Length	Reaction Latency	Duration	Traffic	Target
Infrastructure Protection	Automatic Detection and Triggering	/32	Secs-Mins	Mins-Hours	Attack	Server
Prefix Squatting Protection	Manual	≤ /24	NA	Months	Scanning	None
Content Blocking	Manual	/32	NA	Weeks-Months	Normal	Server

correlated anomalous traffic peak directly before the first RTBH during an attack event. Since 75% of DDoS attacks are volumetric attacks [31], we expect to see a change in the port distribution, i.e., more traffic from amplification candidate protocols often used in DDoS attacks such as DNS, NTP, or memcached. The average duration of DDoS attacks was 218 minutes by the end of the year 2018 [32]. We also assume that servers are a frequent target of DoS attacks, but also attacks to clients have been observed before [3]. Attacking business-critical, often used servers allows the attacker to exert pressure and blackmail the victims for financial gain. Profit margins of a DDoS attack can reach up to 95% [23]. Consequently, we should be able to observe legitimate, regular traffic patterns and compare them with attack traffic, which allows quantifying collateral damage of RTBH as a DDoS mitigation approach.

Based on measurements with an Internet telescope, related work shows that [16] RTBH is usually triggered automatically after a short reaction time. We expect to see this behavior also at our vantage point. RTBH triggered by DoS mitigation mechanism should be rather short, optimally only for the duration of the attack.

In summary, we anticipate the following order of actions for this use case: First, the attack event takes place in the form of a DDoS attack. This distributed attack utilizes multiple attack vectors and attacks either state (e.g. TCP Syn attack) or capacity (UDP-Amplification) of its victim. The attack starts with the increase of unwanted traffic and ends with its disappearance. Second, two parties might react to the attack event. Either the victim itself announces a RTBH or one of its upstream providers, whose links are a collateral damage of the attack. Since all traffic is dropped, the victim is blinded about the progression of the attack. Hence RTBHs will be withdrawn to test for attack traffic and then re-announced. Not only bogus traffic is dropped but also legitimate flows, which is the collateral damage of the mitigation mechanism. We expect to see different traffic properties for the legitimate and attack traffic.

2.3 Prefix Squatting Protection

The increasing scarcity of freely available IPv4 address space and its importance not only for legitimate businesses but also for spammers alike increases the pressure on unused IPv4 address space. Prefix hijacking is a well known phenomenon where IP prefixes are taken over by third parties on the Internet, either erroneously or with malicious intent [27, 41]. Mitigation techniques such as RPKI exist, but are still not sufficiently deployed to completely prevent prefix hijacking [39]. IP prefix squatting is a variant of prefix hijacking, where third parties take over address space that is assigned to another AS but not announced from this legitimate origin [47]. These prefixes are easier to hijack because there is no competitive announcement [2, 22, 36, 41].

One common mitigation technique for prefix squatting is to announce the assigned address space. To ensure the address space is not used at the same time, the same prefix is announced as an RTBH.

Prefix squatting is used in practice, *e.g.,* to send email spam from valid address space and to prevent backtracking [2, 27], or for internal infrastructure addressing in case of address shortage [18]. Considering the severe negative consequences of prefix squatting and the low effort to mitigate, we expect to see applications of this use case in the wild. In fact, we find very few incidents that may refer to RTBH to protect against prefix squatting.

2.4 Content Blocking

Applying RTBH to block clients from accessing content occurs rarely but is possible. Giotsas *et al.* [10] found that attackers (*e.g.,* port scanners, vulnerability scanners) and not victims have been blocked by network operators to prevent access to server content.

Another motivation for the deployment of BGP blackholing is censorship. RTBH can be used to block traffic towards an IP address hosting undesirable content. Compared to access control lists, RTBH reduces operational burdens as it simplifies the maintenance of blacklists [11]. Instead of configuring ACLs on every router separately, a single router maintains the master file and signals the blackhole routes to the peers via BGP. This is specifically beneficial in scenarios that require frequent and rapid changes. We consulted several network operators whether this case has been observed in real-world. Even though the answer was negative, we include this use case for completeness.

In both scenarios, RTBH is characterized by midterm, stable RTBHs routes, triggered by few BGP updates. In particular, *blocked traffic does not reflect typical DDoS traffic patterns.*

2.5 Expected Characteristics

All three RTBH use cases (infrastructure protection, squatting protection, and content blocking) are expected to exhibit different characteristics in terms of BGP signaling and data traffic. Both content blocking and squatting protection are expected to show long-term and stable RTBH routes without attack traffic. In terms of prefix lengths, however, they should differ. Content blocking is expected to use very specific prefixes, *e.g.,* `/32` to filter the addresses of content hosts. Squatting protection, on the other hand, is

expected to predominantly cover ≤ /24 prefixes. RTBH usage for infrastructure protection is expected to show very specific prefixes, similar to content block, but should exhibit DDoS attack traffic during or shortly before the RTBH. We summarize our observations in Table 1. It is worth noting that the classification is not strictly exclusive but indicates common tendencies.

3 DATA CORPUS

This analysis is based on data from a large European IXP that offers remotely triggered blackholing as a service to its members. Our data sets contain three months of passive control as well as data plane measurements. Since more than 95% of the traffic and more than 98% of RTBH events at this IXP are from IPv4, we focus this work on IPv4, only.

We now present our data sets in more detail including the different data sources, potential challenges when aligning different sources, and a brief overview about blackholing activity at our vantage point.

3.1 Control and Data Plane Data Sources

We measure data on the control plane (*i.e.,* BGP) to identify remotely triggered blackholing. For a better understanding of the RTBH impact on the blackholed IP prefixes, we capture traces on the data plane (*i.e.,* IPFIX). All measurements are consistently taken from September 26, 2018 until January 11, 2019. We had to exclude few hours and December 6 due to infrastructure maintenance.

Control Plane. An AS initiates (or terminates) RTBH at the IXP by sending BGP update messages with a specific BGP community [17] to the public IXP route server, which distributes this information further to either all of its peers or to a subset. We collect these messages and gain the following information: (*i*) when the blackholing should start and stop, (*ii*) which AS triggered RTBH, (*iii*) which ASes should send data to the blackhole, and (*iv*) the origin AS of the RTBH prefix. The time resolution of the collected, RTBH-related BGP messages relies on the NTP protocol for synchronization and, therefore, is expected to be accurate at a level of 10ms [30].

Note that RTBHs established in bilateral (private) peering is out of scope of this paper.

Data Plane. We collect IPFIX packet samples (1 out of 10,000 packets) of incoming traffic from peers at all member-facing ports of all network devices at the edge of the IXP switching fabric. On average, we sample 70,000 packets-per-seconds. From the collected packets, we extract the packet sizes, source and destination MAC addresses, destination IP addresses, source and destination transport ports. Based on this data, we can attribute 590 million packets as originated from or addressed to any of the blackholed IP prefixes. To identify the ASes that exchange the packets at the IXP, we map source and destination MAC addresses of the sampled packets to the router interface addresses of the ASes connected to the IXP switching fabric. This collection includes 47,000 IPFIX flows received from internal system of the IXP as a source or destination device, *i.e.,* 0.01% of the total number of flows. The internal traffic is removed from the data set before further processing. Thereby, we have full, 1 in 10000 sampled visibility of all member traffic coming into the IXP switching fabric. This data set is used for analyzing both forwarded traffic and dropped, blackholed traffic.

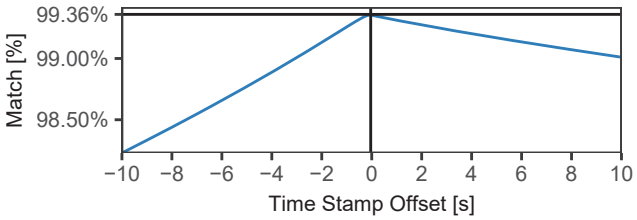

Figure 2: Maximum likelyhood estimate for time offset between control and data plane sources.

Identifying Dropped Traffic. The dropping of blackholed traffic at this IXP is implemented with the help of a unique (blackhole) MAC address that does not forward data. By announcing a special next hop via BGP, which in turn maps to this MAC address, we redirect packets to the blackhole. Consequently, the data is dropped, and we can mark any sampled packet with destination to the blackhole MAC as dropped traffic.

Using the sampled data of dropped packets in correlation with our control plane measurements, we calculate the amount of dropped traffic triggered via the route server. We find that on average, 95% of the dropped bytes are controlled by RTBH signaled via the route server and therefore represent the majority of the observed traffic. The remaining 5% belong to traffic that was dropped because of other RTBH sources.

Accuracy of Timestamps. All measurement devices synchronize their system time using NTP in the local subnet, which allow for a time series analysis between both data sets. Deviations, however, are still possible and need careful verification. To quantify errors, we measure which share of the sampled packets was dropped because of blackhole announcements visible in the recorded BGP data and which share was not dropped. Based on the timestamps from the control and data plane, we apply a maximum likelihood approach to estimate the time offset between both data sets.

During the measurement period, ≈ 50 M packets addressed to RTBHs were dropped by the blackholing service. The offset between the control and the data plane is depicted in Figure 2. The maximum overlap is 99.36% for an offset of $-0.04\,s$, showing that both data sources are sufficiently consistent in time.

Note that control plane data is needed for the subsequent analysis as we are also interested in events where BGP announcements signal RTBH but the receiving ASes still forwards data (*i.e.,* does not select the announced RTBH prefix as best route).

3.2 RTBH Load

Figure 3 provides an overview on the load of the RTBH signaling. During the measurement period, 830 member ASes have been connected on average to the IXP peering platform. 78 of these peers announced 1,107 RTBHs for 170 origin ASes at any given minute in the observation period. At most 1,400 RTBH prefixes were active during the same minute, which is less than two prefixes per connected peer. The number of RTBH-related BGP messages stays

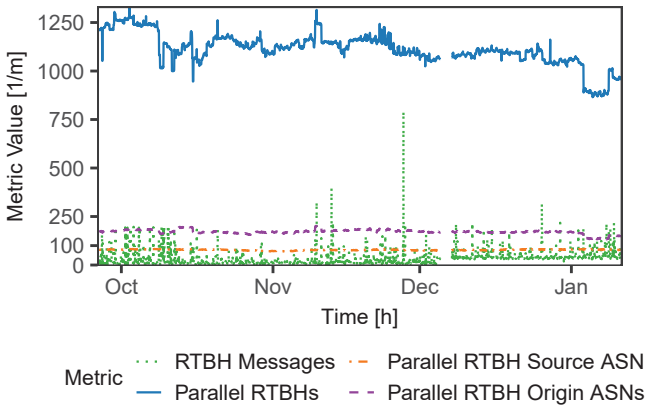

Figure 3: Number of active parallel RTBH over time.

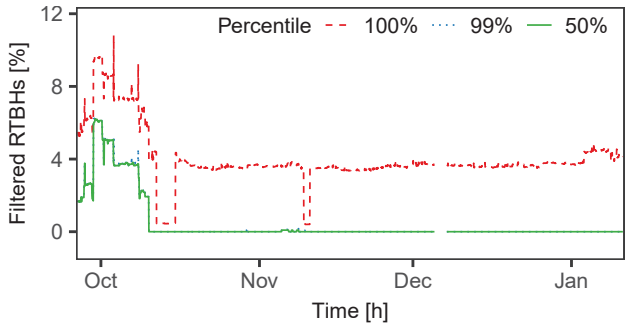

Figure 4: Percentage of all announced blackholes at a given time that are filtered and are not visible to 100/99/50 percentiles of peers on the peering platform. 99% and 50% quantiles overlap such that only the 50% quantiles are visible in large parts of the figure.

below 500 messages with a few spikes of up to 600 and one spike up to 793 messages per minute or less than 14 messages per second.

These numbers illustrate nicely that RTBH adds negligible overhead on the control plane in terms of memory and processing. This resource efficiency might explain the popularity of using RTBH to protect the Internet infrastructure.

4 ACCEPTANCE OF RTBH FEATURES

The efficacy of RTBH-based filtering relies on both receiving related BGP announcements from the route server *and* accepting the received routes as best paths. In this section, we answer the two questions: Do network operators try to reduce the negative impact of RTBH? Do network operators accept RTBH announcements to filter traffic?

4.1 Using Targeted Blackhole Routes

The RTBH service at our vantage point allows network operators to instruct the route server to selectively announce RTBHs to specific ASes on the peering platform, which reduces collateral damage. Using BGP communities, the victim AS can select to which peers its RTBH announcement will be forwarded by the route server. Thereby, unfiltered communication continues with unaffected neighbors.

It should be preferential for an operator to affect only the ASes transporting malicious traffic by RTBH. We investigate this hypothesis by analyzing the BGP communities which are collected in our control plane data set. Thereby, we are able to obtain the specific view of every BGP peer on the set of blackholed prefixes at every point in time throughout the measurement period.

Figure 4 shows the percentage of all announced blackholes at a given point in time that are filtered to not be visible at all peers. The quantiles indicate which part of the announced blackholes are filtered and, therefore, not visible to all (100%), 99%, and the median (50%) of the connected peers. Significant deviations of parallel RTBHs are visible during some weeks at the beginning of October 2018. At this time, the median of the peers saw up to 6.2% fewer RTBHs than the route server and a single peer even 10.8% fewer. After mid October, however, the median and 99% percentiles of the peers dropped down to at most 0.2% fewer RTBHs compared to the full visibility at the route server. The peer with the fewest received RTBHs saw only a minus of up to 4.9% parallel RTBHs. Based on these findings we conclude that selective filtering and announcements are the exception and commonly not used to reduce the collateral damage for targets of DDoS attacks.

4.2 Accepting Blackhole Routes

Any BGP peer that does not accept a blackhole route from the route server will continue to forward the traffic that was intended to be filtered. Acceptance of this route is beyond the control of the triggering AS, but subject to local BGP policies of the receiving peer. Using the RTBH visibility information derived in Section 4.1, we calculate the fraction of data that a router transmits even though it received a blackhole announcement.

Figure 5 depicts the amount of traffic dropped for all active blackhole prefixes relative to the overall amount of traffic for those prefixes during the blackhole distinguished by prefix length. The opacities of the colors visualize the RTBH traffic share of the respective prefix lengths compared to the overall blackhole traffic. For example, 99.9% (highest opacity) of the overall RTBH traffic was sent to /32 prefixes and < 0.01% (lowest opacity) of the traffic to

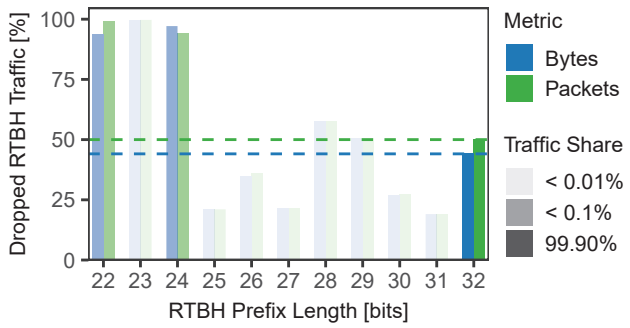

Figure 5: Observed shares of dropped traffic by RTBH prefix lengths; dashed lines denote averages. The traffic shares are visualized as opacities of the bars.

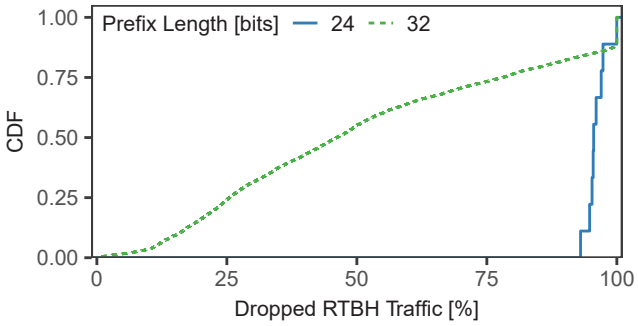

Figure 6: Distribution of dropped RTBH traffic shares for selected prefix lengths.

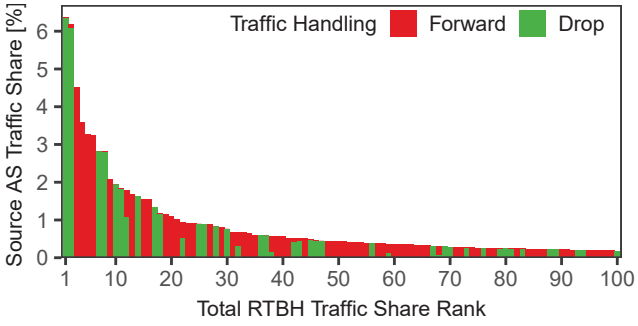

Figure 7: Reaction of top 100 source ASes by traffic share to /32 RTBHs.

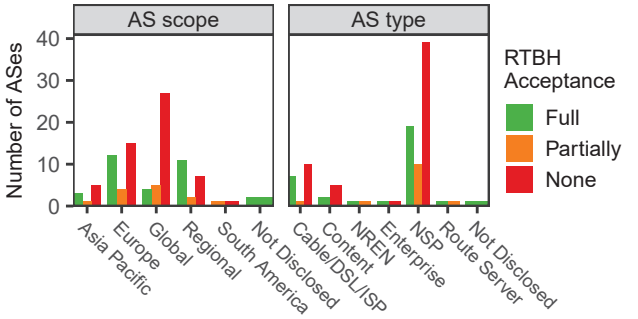

Figure 8: The PeeringDB organization types of the top 100 source ASes by traffic share sent to /32 RTBHs.

/25 or /26 prefixes. The dashed lines show the average drop rates of RTBH announcements considering all RTBH prefix lengths.

It is clearly visible that the vast majority of traffic for blackholing corresponds to /32 prefixes. To our surprise, however, only 50% of the packets (or 44% of bytes) are filtered, *i.e.*, more than half of the traffic continues flowing to the victims. In contrast, blackhole routes to less specific prefixes (/22, /23, and /24) are accepted as best paths in 93% – 99% of the cases. Those prefix lengths are common in BGP announcements in general [43]. Considering that more specific prefixes (/25 - /31) exhibit a behavior similar to /32 in terms of the dropped rate, we assume incorrectly configured BGP policies because accepting (RTBH) prefixes longer than /24 bits requires to change the common BGP configuration and to whitelist such announcements.

To better understand the varying acceptance for different prefix lengths, we investigate the behavior of /24 and /32 prefixes in more detail. Figure 6 shows the CDF of the observed drop rate for these two prefix lengths. The drop rate of traffic to /24 RTBH prefixes varies between 82% and 100% with a median of 97%, making /24 blackholes a fairly predictable configuration to successfully mitigate unwanted traffic.

For /32 prefixes, the blackhole traffic drop share ranges between almost zero and 100%, with 30% for the first quartile, 53% for the median, and 88% for the third quartile. This wide distribution results in a high uncertainty regarding the expected effectiveness when announcing an /32 RTBH. In the median case, the unwanted traffic will be reduced to at least half, but in some cases an RTBH announcement will cause no data reduction at all. Triggering RTBH for single hosts (/32 prefixes) is often very appropriate, but may lead to a rather unpredictable reaction in reducing unwanted traffic.

To characterize the AS peers further that ignore /32 announcements and cause low drop rates, we investigate the top 100 source ASes that contribute most of the traffic volume to /32 blackholes. Figure 7 shows the relative amount of dropped and forwarded traffic by these ASes that all together account for over 85% of the total traffic to RTBHs, many of which are heavy hitters in RTBH scenarios. Only 32 of these ASes drop more than 99% of the traffic to RTBHs. 55 of the top source ASes forward only less than 1% of the traffic to the blackhole route. Interestingly, 13 ASes exhibit an inconsistent behavior as they send significant parts of the traffic to the blackhole route and forward other parts of the traffic to the victim AS.

For a deeper dive, Figure 8 groups the top 100 ASes by their AS types and scopes based on PeeringDB data. Most ASes that do not (or partially) accept blackhole routes are network service providers (NSPs), which comes as a surprise. We expected these companies to be well-prepared for complex BGP configuration tasks. One reason for the contrary may be that global NSPs deploy alternate measures of DDoS mitigation, outside the public peering ecosystem.

5 EVIDENCE OF DDOS ATTACKS

The default use case of RTBH is considered DDoS protection. In this section, we explore this common assumption by correlating events at the data and control plane. This analysis requires a careful modeling of common DDoS and mitigation patterns to differentiate the signals at the control and the data plane.

5.1 Preparatory Steps

Blackholes for infrastructure protection are announced and withdrawn repeatedly to check whether the attack event is still ongoing (see Figure 9). In practice, as we have shown in Section 4, some traffic arrives even when a blackhole is active and thus might serve as indicator for an ongoing attack. This remaining traffic, however, is a highly unreliable source of status information due to the high variance in actual drop rates. That is why we still see frequent re-announcement patterns.

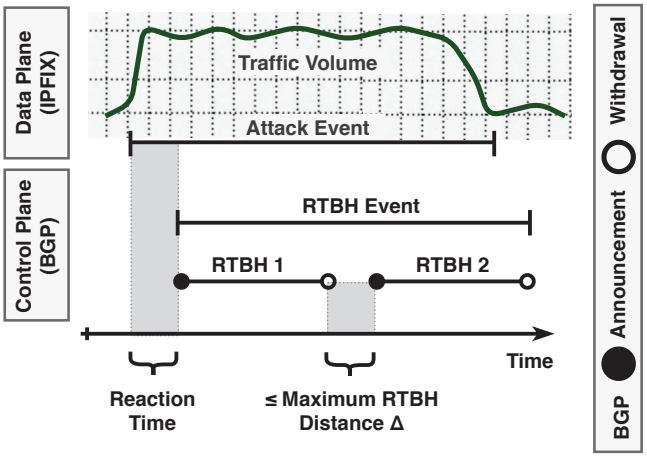

Figure 9: Attack and RTBH events: A sequence of re-announced RTBHs.

To consider multiple RTBH announcements that target at the same *attack event*, we group on-off update patterns into a single *RTBH event*. Each RTBH event reflects the mitigation process after the attack was detected. Small gaps between multiple RTBH announcements that belong to the same attack event are likely to show attack traffic as well. To prevent the misclassification of traffic, we include traffic during these gaps into RTBH events. The challenge is to find an appropriate time threshold Δ between consecutive RTBH announcements, which distinguishes RTBH announcements that belong to the same or another RTBH event.

For each blackhole update bh_i of a single RTBH event, the following applies with respect to the observed timing between BGP withdrawals and announcements:

$$|\, bh_i[\text{withdraw}] - bh_{i+1}[\text{announce}] \,| \leq \Delta$$

Now, we need to find an appropriate merge threshold Δ. For this, we consecutively increase Δ and inspect the amount of blackhole events, relatively to the overall number of RTBH announcements (see Figure 10).

The last significant effect is visible up to about Δ = 10 minutes. Furthermore, a 10 minute Δ is consistent with the delay found between the detection of DDoS traffic and the triggering of a blackhole [16]. Therefore, even if the blackhole originator mistakenly disables a blackhole while an attack event is still ongoing, a newly triggered blackhole would be part of the correct, preceding blackholing event. For this Δ, 400k blackhole announcements are grouped into only 34k RTBH events, which is a reduction to 8.5%. We highlight the lower bound Δ = ∞ (red dashed line), for which the number of RTBH events equals the number of unique blackholed prefixes.

Fixing the merge interval to the reasonable threshold of Δ = 10 minutes, we can now use the aggregated RTBH events to examine the traffic before and during RTBH events.

5.2 Visibility of Pre-RTBH Events

Assuming that most RTBH events are triggered by volumetric DDoS attacks, we search for traffic anomalies during the 72 hours before

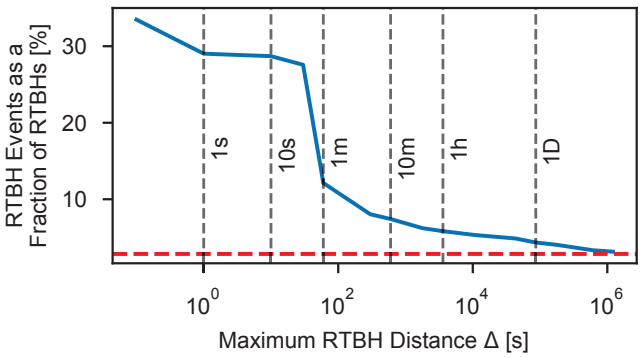

Figure 10: Fraction of blackholing events in all RTBH announcements.

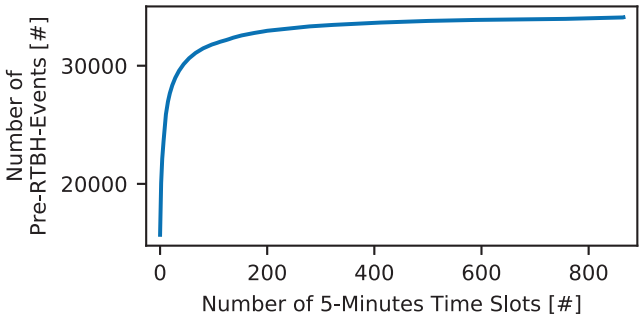

Figure 11: Cumulative number of time slots contributing traffic samples within 72 hours before RTBH started.

the first RTBH announcement. We refer to this time range as the pre-RTBH event. If a pre-RTBH event contains an anomaly, the corresponding RTBH event is said to have a preceding anomaly.

First, we identify all pre-RTBH events that include at least one sampled packet and thus may give additional insight into the traffic behavior. We aggregate into five minutes slots and show the cumulative contributions in Figure 11. Surprisingly, traffic appears for only 18k of the total 34k pre-RTBH events. This means that 46% of all pre-RTBH events did not exhibit a sufficient amount of packets to be sampled, even though our vantage point is one of the largest IXPs. For these cases, data plane monitoring cannot explain the root cause of the RTBH events. 13k of pre-RTBH events exhibit data for at most 24 time slots during a total of 2 hours (see Figure 11). This indicates very sparse data. Manual inspection shows that those pre-RTBH events represent incidents where unusually high traffic peaks are visible shortly before the first RTBH announcement. This motivates further investigation, which we continue in the next section.

5.3 Classification of Pre-RTBH Events

We want to automatically describe and classify the traffic behavior for the pre-RTBH events. For this, we observe five traffic features: (*i*) number of packets, (*ii*) number of flows, (*iii*) number of unique source IP addresses, (*iv*) number of unique destination ports, and (*v*) number of non-TCP flows.

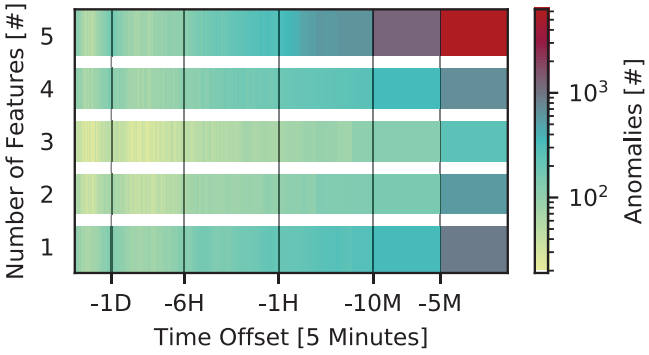

Figure 12: Level and Time Offset of Traffic Anomalies during pre-RTBH events.

As a straightforward indicator, we use Exponentially Weighted Moving Average (EWMA), a simple sliding window mechanism to detect unusual traffic peaks. For each detection, we consider a 24 hours window, which shifts every five minutes and spans 288 time slots. The most recent values have the highest weight, the oldest the smallest weight. We use the same notation as common data analysis tools [33]. The decay parameter α and the weight w are calculated as follows:

$$\alpha = 2/(s+1), \text{with } s = 288$$
$$w_i = (1-\alpha)^i$$

Then, the weighted moving average is defined as

$$y_t = \frac{\Sigma_{i=0}^{t} w_i x_{t-i}}{\Sigma_{i=0}^{t} w_i}$$

Please note that we require a full window for an anomaly detection. This means that no anomaly can be found during the first 24 hours. We perform an EWMA anomaly detection independently for each feature. Values are tagged as anomalous when they exceed the moving average by $2.5 \cdot \text{SD}$ (standard deviation). Then, we count the number of features that have an anomalous traffic peak for each time slot. We refer to this as the anomaly level.

Figure 12 shows the distribution of all anomalies by level and time offset relatively to the RTBH event start. There is a clear trend for the time-lag between anomalies and RTBH events in that most anomalies occur up to ten minutes before the first RTBH announcement. This short reaction time indicates automatic DDoS mitigation tools. Usually, all five features show anomalous behavior shortly before the blackhole. We also find multiple cases in which an anomaly was found only for one of the five features. This emphasizes the importance of a multi-sided traffic analysis to detect individual anomalies.

Based on these results, we now are able to classify pre-RTBH events into three classes: Pre-RTBH events (*i*) without sampled traffic, (*ii*) with sampled traffic, but no anomaly before the RTBH event, and (*iii*) with sampled traffic and at least one anomaly before the RTBH event. The first class has been quantified in the previous section.

We find 9k pre-RTBH events (27%) with an anomaly up to 10 minutes before the initial RTBH announcement. Also, we find only

Table 2: Class Distribution of Pre-RTBH events.

Pre-RTBH Event Class		% Events
Data	Anomaly ≤ 10 min	
✘	–	46%
✓	✘	27%
✓	✓	27%

11k pre-RTBH events (33%) with an anomaly up to 1 hour before the initial RTBH announcement. This means that only one third of all RTBH events are triggered by volumetric traffic changes. This finding deviates significantly from the original intention of RTBH as a tool of DDoS mitigation. We summarize these results in Table 2.

Relevance of Anomalies. We identify DDoS anomalies by volume, no matter whether they are the result of spoofed or unspoofed, direct or reflected attacks. Our approach detects sudden peaks but is not able to detect long-tailed DDoS attacks such as Slowloris [4]. Long-tailed DDoS attacks, however, do not produce large traffic volumes and are not expected in the context of RTBHs.

The median DDoS attack size in mid 2018 was 1,287 Mbps [25]. Dividing by a MTU of 1,500 Bytes, this corresponds up to 100k packets per second on the IXP switch fabric. Due to the large number of packets even for medium size attacks, we expect the observed anomalies to be visible also in our sampled traffic data.

It is challenging—if not impossible—to verify our methodology based on ground truth data because most companies are hesitant to reveal such information. The common attempt of correlating the results with public documentation of DDoS attacks is not necessarily helpful, either, due to the deployment of other mitigation tools. During our measurement period, for example, Imperva reported one of the largest attacks ever observed [42] but their mitigation portfolio contains only scrubbing centers and DNS diversions and not RTBH.

Nevertheless, we tried to verify the largest attacks visible in our data set. We were successful in some of the cases, even though many companies retain information on attacks from the public, as they do not want to admit reachability problems. As one example, an online shop that experienced the fifth largest RTBH event by attack volume, confirms the attack and time in a public announcement [38]. In this case, we identified an active RTBH event (with a preceding anomaly) that lasted for more than 7 hours.

To ensure the correct detection of traffic spikes, we performed multiple consistency checks and manual inspections. Fortunately, a clear trend becomes apparent, which justifies our simple detection methodology. Either we do not observe any traffic changes at all or very significant bursts. To further substantiate this observation, we now analyze the relevance of the anomalies compared to the average traffic behavior 72 hours before the RTBH event begins. Since most anomalies occur ≤ 5 minutes prior to the RTBH event, we focus on this time slot.

For every traffic feature, we calculate the relative rise during the last five minutes prior to the RTBH event, which we refer to as *Anomaly Amplification Factor*. This factor is depicted in Figure 13. The time slots covering ≤ 5 minutes often do not contain traffic, either because the entire pre-RTBH event does not contain any

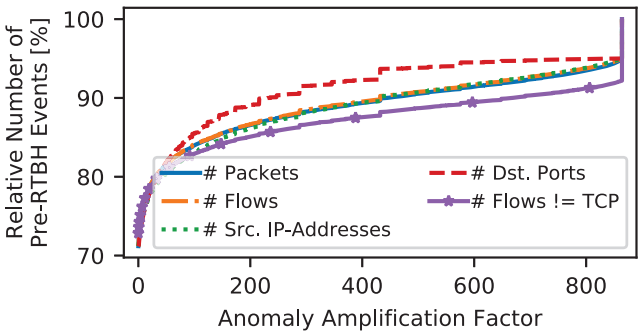

Figure 13: Last time slot compared to the mean of the respective pre-RTBH event.

data at all, or because packets are only seen in other time slots. Nevertheless, if packets were sampled during the last five minutes, large multiples of up to 800 can be observed. In 15% of the cases this slot shows the maximum value of the entire time range.

These results indicate very strong changes in traffic patterns that occur with the anomalous behavior. It does not require any fine-tuning. Instead, we tested extreme configurations such as thresholds of $10 \cdot SD$ (instead of 2.5) with very stable results.

5.4 Classification of RTBH Events

We now inspect the traffic during the RTBH events. Even though we sample packets at an Internet exchange point of very large data volume, the sampling does not necessarily capture packets during an RTBH event. To gain insight into this general measurement challenge, we first classify the events according to traffic visible on the data plane. Then, we correlate the visible traffic with commonly misused services.

Overall, the sampling captured packets for only 29% of all RTBH events, albeit we applied a high sampling rate of 1 out of 10,000 packets. More than half of the RTBH events that feature captured data have also a preceding anomaly within 10 minutes. These incidents account for 18% of all RTBH events. Interestingly, one third of the RTBH events with a preceding anomaly have no traffic during the RTBH event. We explain this by (*i*) very short-lived DDoS attacks and (*ii*) other mitigation points on the Internet that drop the attack traffic before reaching our vantage point (*e.g.*, scrubbing [16]).

We now analyze the network service misused to generate attack traffic. We expect to observe attack traffic during RTBH events with a preceding anomaly. This is why we identify the protocol distribution for each RTBH event for which a (preceding) anomaly was detected *and* the monitoring system sampled traffic. We find that UDP is the most prevalent transport protocol in this context: (*i*) 99.5% UDP, (*ii*) 0.3% TCP, (*iii*) 0.1% ICMP, (*iv*) 0.1% other. This protocol distribution differs significantly from the normal traffic mix at IXPs [1, 6].

For the dominant UDP traffic, we check whether the RTBH events relate to attacks based on common UDP amplification protocols. To prevent biased results due to outliers, the RTBH traffic analysis is conducted on a per event basis. It is worth noting that the analysis relies on transport ports because the application payload is not available for privacy reasons. We find that the majority of packets

Table 3: Different UDP amplification protocols* per RTBH event that shows data and preceding anomaly.

Different protocols* [#]	0	1	2	3	4	5
Events [%]	6	40	45	8.3	0.6	0.1

*Considering the following known amplification protocols/UDP ports: QOTD/17, CharGEN/19, DNS/53, TFTP/69, NTP/123, NetBIOS/138 SNMPv2/161, LDAP/389, RIPv1/520, SSDP/1900, Game/3659 Game/3478, SIP/5060, BitTorrent/6881, Memcache/11211 Game/27005, Game/28960, Fragmentation/–.

can be assigned to one or two amplification protocols during RTBH activity (see Table 3). The most common amplifying protocols per event are cLDAP, NTP, and DNS, all of them significantly misused for amplification attacks [19].

Based on our analysis of the protocol mix, we are able to investigate the potentials of fine-grained filtering in the next section. For such an analysis, neither the amount of exchanged packets nor bytes is important. The overall trend of these measures in the context of blackholing is described in related work [6].

5.5 Potentials of Fine-Grained Filtering

Since most events show traffic patterns of well-known attacks, we investigate the impact of fine-grained filtering to prevent collateral damage. For each RTBH event with an anomaly and available traffic, we emulate the filtering of UDP amplification packets. Figure 14 shows the relative amount of RTBH events where filtering of a specific ratio of amplification packets is possible. Fortunately, 90% of the RTBH events could be supported completely by dropping common UDP amplification traffic based on an a priori known port list. Such fine-grained filtering would prevent collateral damage in a lightweight fashion. The remaining 10% require further investigation and are more difficult to mitigate. We observe attacks on random ports, increasing port numbers, and the use of multiple transport layer protocols.

So far, we did not investigate the source networks of the attack traffic in detail. Since we pre-dominantly observe UDP reflection attacks, most source IP addresses are not spoofed but sent from reflectors to victims. This allows us to determine the *origin AS* of

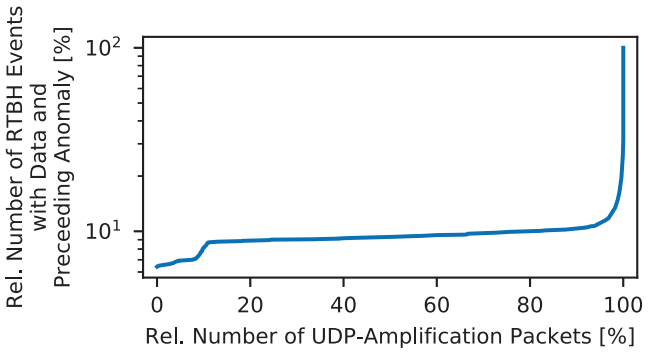

Figure 14: Relative amount of dropped packets per event if filtered by known UDP amplification traces.

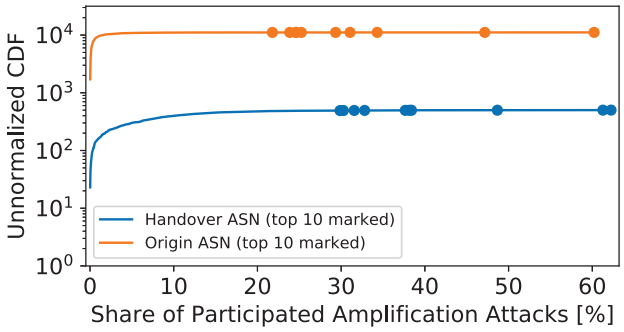

Figure 15: Cumulative number of ASes that participated in a relative amount of UDP amplification attacks. Shares of top 10 handover and origin ASes are highlighted.

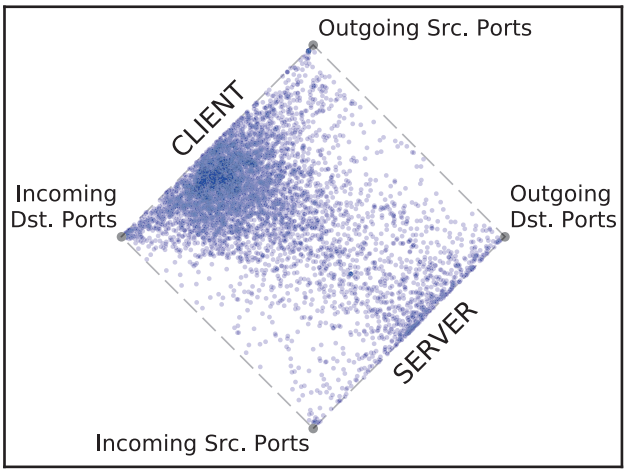

Figure 16: Port Distribution of IP addresses outside of pre-RTBH events.

the attack traffic, *i.e.,* the AS hosting the amplifier. Moreover, we are able to determine the *handover AS*, *i.e.,* the ingress AS at the IXP switch fabric. As this mapping is based on MAC addresses of routers at the IXP, it is also not susceptible to spoofing.

Figure 15 displays the CDF for the share of UDP amplification attacks in which the handover AS and the origin ASes have been the source of an attack. Overall, we observed 501 handover ASes (55% of all IXP members) and 11,124 origin ASes (17% of all advertised ASes) that participate in attacks events. The majority of handover ASes at the IXP do not participate in more than 10% of the events and most origin ASes do not participate in more than 3% of the events. Complementary, we also find a few ASes that have been involved in 20%-60% of the attacks. We highlight the 10 last discrete steps in the CDF, which mark the top 10 ASes in each category. The top-ranked origin AS (60% of events) and the handover AS (62% of events) are the same AS. Although participating in so many events, this origin AS is only responsible for 6% of the total attack traffic. On average, we observe 1,086 amplifiers during an attack and traffic from 30 handover ASes or 73 origin ASes. Our results indicate a highly distributed usage of amplifiers, which makes fine-grained blacklisting based on the attack source very difficult.

6 INVESTIGATING COLLATERAL DAMAGE OF RTBH

We have analyzed the blackholed traffic without inspecting the legitimate traffic. We will now try to identify legitimate traffic based on reoccurring traffic patterns outside of the RTBH events. Such information could be used to implement whitelists during an attack, and to approximate the collateral damage during RTBH events.

6.1 Port Distribution per Host

Since we observed traffic anomalies before RTBH events (see Section 5.3), we prepend a 10 minutes reaction time to each of these RTBH activities. Traffic during this reaction time is not classified as legitimate. We select hosts (identified by an IP address) with incoming and outgoing traffic on at least 20 different days, which is a conservative lower bound of samples to identify legitimate traffic. Only 30% of blackholed IP addresses meet this criteria. To verify our assumption that servers are a common DDoS victim and hence tend to be blackholed, we first need to distinguish server hosts from client hosts. Therefore, we inspect four features:

(1) # of unique source ports in incoming traffic
(2) # of unique source ports in outgoing traffic
(3) # of unique destination ports in incoming traffic
(4) # of unique destination ports in outgoing traffic

We expect the following behavior based on a common client-server scenario. A server should receive traffic on few dedicated listening ports. In contrast, clients use random source ports to initiate communication. The server will thus receive traffic from many different ports and reply to these many ports from its stable ports.

We use a RadViz projection [13] to visualize our results, see Figure 16. RadViz visualizes multi-variate data by projecting an N-dimensional data set into a 2D space. Features are represented by anchor points equally spaced around the perimeter of a unit circle. Each data point is attached to all anchors by a spring, the stiffness of which is proportional to the numerical value of that feature. The values are normalized by the maximum number of values each feature can attain. Data points are closer to the anchors for which they have higher values than for the others.

In our case, each data point represents a host, the features represent the ports, and the normalization factor is derived by the maximum port number (*i.e.,* 1/65535). Client hosts will be pulled by an anchor that represents high diversity in the number of unique destination ports in incoming traffic (or high diversity in the number of source ports in outgoing traffic). On the other side, server hosts that send traffic to clients will be pulled by an anchor that represents high diversity for source ports in incoming traffic (or high diversity for destination ports in outgoing traffic).

We observe more IP addresses that show traffic patterns of clients (see Figure 16). To our surprise these nodes are protected by RTBH.

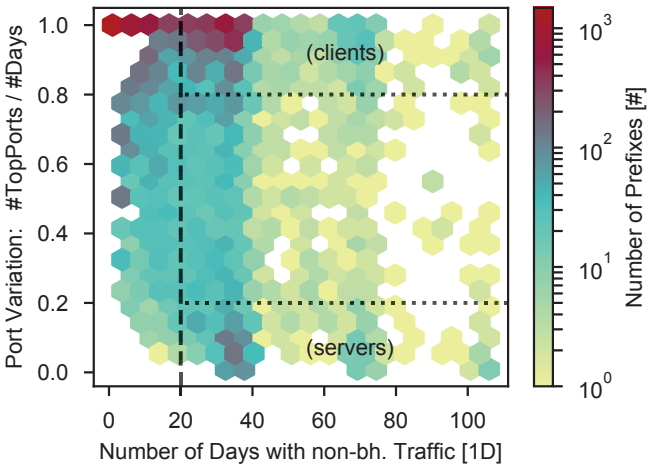

Figure 17: Top port variation and classification of IP addresses for traffic outside of RTBH events.

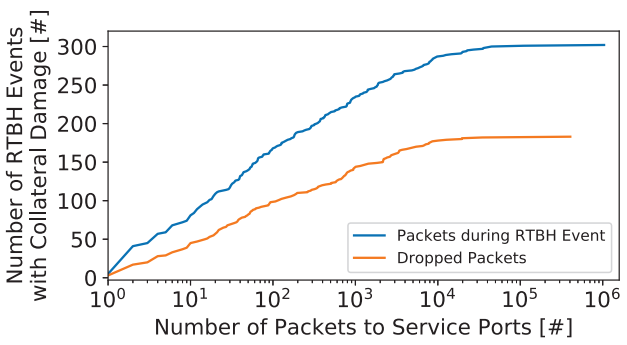

Figure 18: Collateral damage during RTBH events for servers. We differentiate by all packets to service ports and actually dropped packets.

6.2 Detecting Stable Traffic Patterns

To get a better understanding of the previous observation, we refine our results by inspecting the incoming traffic in more detail. This analysis is particularly challenging as client traffic is highly variable, which makes the detection of normal traffic patterns difficult.

For each destination IP address, we determine the number of days with incoming traffic and for each day the most utilized destination port, which we call *top port*. Note that we differentiate between protocols, so each port is identified by a protocol-port-tuple, *e.g.*, (TCP, 80). Based on this, we compute the port variation, which is the ratio between the number of top ports and days with traffic. Consequently, a port variation of 1 means that we have observed a different top port on each day. A port variation close to 0 indicates very stable top ports, which resembles the behavior of frequently used servers with well-known applications. We show our results in Figure 17. We use the port variation to classify hosts as clients or servers. Again, we require at least 20 days of captured packets. We find over 4,000 clients and 1,000 stable servers.

To gain confidence in our results, we map each client and server IP address to its origin AS. Then, we retrieve the AS type from Peering DB [35], see Table 4. The most common AS type for clients is "Cable/DSL/ISP" (60%), for servers it is "Content" (34%). This means we found over 2,000 hosts with traffic patterns resembling clients that are actually located in ISP networks and have been targeted by DDoS attacks. DDoS attacks on clients have been reported before [3, 44, 45]. These attacks occur mainly due to disputes in online gaming and to manipulate e-sport matches [44]. Nevertheless, we are surprised how pronounced this shows up in our data set, in particular in comparison to the identified number of traditional servers.

Table 4: ASN types for detected server and client IP addresses based on Peering DB.

Type	Clients	Server
# Hosts	4057	1036
Content	2%	**34%**
Cable/DSL/ISP	**60%**	14%
NSP	14%	13%
Enterprise	1%	1%
Unknown	23%	38%

6.3 Towards Quantifying Collateral Damage

The identification of servers with stable top ports allows us to present a preliminary assessment of the collateral damage during RTBH events. Note, that clients have a different top port for almost every day of activity. This makes a description of legitimate traffic patterns very difficult. In contrast, the detected servers have only a small list of frequently addressed top ports, which indicates legitimate traffic patterns.

For each detected server, we quantify the number of sampled packets sent to the identified top ports during RTBH events. Overall, we find 300 RTBH events with traffic including collateral damage for our 1000 detected servers. The (unnormalized) CDF for the number of packets to top ports is shown in Figure 18. We differentiate between all packets sent to top ports during an active RTBH event, *i.e.*, packets that should have been dropped, and those that were actually dropped. We deliberately decided not to quantify collateral damage as a relative traffic share. Expressing collateral damage in percent yields very small shares which only point towards large attack volumes, which are expected during DDoS events. Hence, in order to quantify the collateral damage, we show the absolute values. We observe a collateral damage of up to 10^6 packets. Note that we cannot differentiate between collateral damage and attack traffic sent to top ports, *i.e.*, application specific attacks. Thus this graph shows the upper limit, worst-case, of collateral damage for the detected servers.

Understanding the Challenges of Future Work. Based on our analysis, we identify the following challenges for the assessment of collateral damage. First, we detected servers and clients as victims of DDoS. Since clients have variable usage patterns and might also receive

dynamic IP addresses from IP address pools, finding stable patterns for these cases is very difficult.

Second, we see two sources of bias in our traffic captures. (*i*) Incoming traffic is biased by scans. End-hosts might receive traffic on ports although no application is listening on that ports. (*ii*) Outgoing traffic is biased by spoofing. Spoofed packets suggest traffic from ports, which the end-host actually never used.

Third, in most cases we have very sparse data outside of RTBH events. This impedes results that are statistically significant. Packet sampling does not only reduce the number of packets visible, but also the level of information. We only see header-data up to the transport layer without the possibility to interpret application payload for a finer service-detection.

Fourth, attack traffic is also present outside of RTBH. We deal with this challenge by inferring RTBH events. However, not all DDoS attack have to trigger a DDoS mitigation. It remains open, whether the RTBH information we collect from the route-server is sufficient for a reliable traffic classification.

Last, the patterns of legitimate traffic might change during a DDoS attack. For example, legitimate clients will send more Syn-requests to a server which is not responding due to being overloaded. This behavior has been also observed for stateless protocols such as DNS over UDP and was termed friendly-fire [29].

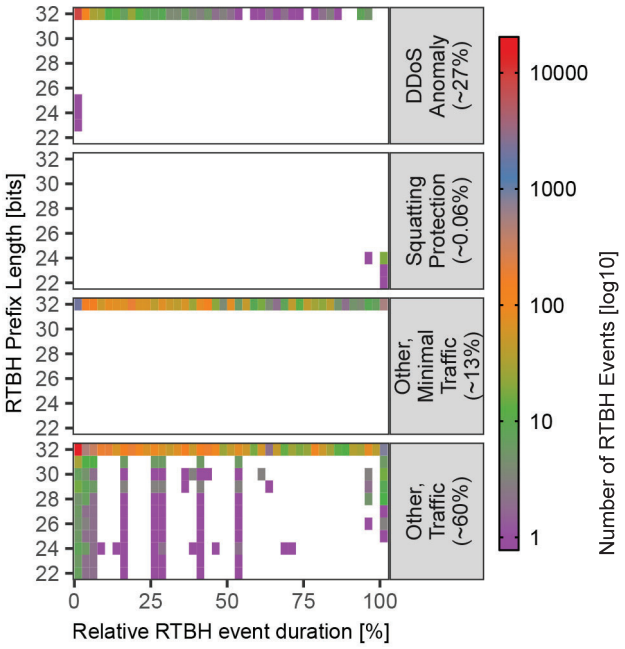

Figure 19: Classification of RTBH events according to different use cases.

7 DISCUSSION OF FINDINGS AND OPERATIONAL PRACTICES

In this paper, we have investigated different perspectives on RTBH from the scientific point of view. The actual usage patterns of RTBH, however, are strongly influenced by practical considerations of network operators. Therefore, we discuss how the findings in this paper can give insights from a practical point of view.

7.1 RTBH Acceptance

Most BGP routers available on the market today support RTBH with small configuration adjustments. The default BGP configurations of virtually all devices, however, do not accept prefixes longer than /24, yet—including blackhole announcements.

Specific configuration settings are required to accept longer prefixes for blackholes only. Our investigation into the RTBH acceptance rates by prefix length show that RTBHs with prefix lengths between /24 and /32 exhibit especially low dropped packet shares. The most likely reason for that is that some operators specifically enable whitelisting of /32 prefixes in their routers, but not the prefix lengths between /25 and /31.

More importantly, our investigations also showed that the number of operators that do not accept /32 blackhole routes is alarming. Surprisingly, this does not only affect small and mid-sized network operators, but some of the largest network operators connected to the IXP. Only 32% of the 100 top traffic source ASes accept host-specific blackhole routes, which are typically used to mitigate DDoS.

The deployment and usage of these incomplete RTBH configurations do not only lead to unpredictable protection against unwanted traffic, but may also shed light on the efforts required to enable RTBH. While low margins and high market pressure explain why many small- and mid-sized operators choose not to invest in these configuration adjustments, the reasons why global network service providers remain unclear. One reason might be that large network service providers use alternative mitigation approaches outside of the IXP ecosystem to handle DDoS attacks.

In any case, missing incentives are likely to play a role in the low acceptance of blackholing routes. The ASes that could gain the most from RTBH are under severe market and cost pressure and often lack the necessary skills to implement blackholing correctly. This can be addressed by additional free advanced training of the IXP community. The ASes that do not see the need to use blackholing, either because they can handle the load inside of their network or because they rely on third-party DDoS protection services, are less willing to invest into infrastructure modification that help other ASes only.

7.2 RTBH Collateral Damage Prevention

RTBH is generally a coarse-granular traffic filtering tool. Unfortunately, even the currently available options to reduce the collateral damage triggered by blackholing are not used. Targeted announcements could be used to specifically drop traffic from neighboring ASes that send attack traffic. As we showed, however, the usage of this feature is minimal in the investigated data. Therefore, we conclude that this feature is virtually ignored.

Furthermore, RTBHs could be announced and withdrawn in a timely manner to filter only attack traffic. A significant part of the blackholes, yet, stay active for a very long time (compare the relative RTBH event durations in Figure 19). For those announcements, we found almost no indication that they relate to the alternative explanation, prefix squatting. We therefore suspect that many of

these blackholes were once manually triggered to prevent a DDoS attack and then have been forgotten.

We therefore conclude that preventing collateral damage caused by RTBH-based DDoS mitigation is not a high priority for users of blackholing today. Rather, RTBH is a simple-to-use tool to prevent DDoS attacks that are threatening the network of an operator. Ensuring appropriate reachability of the victims of the DDoS seems to play a minor role in these considerations. The results of these investigations are rather disillusioning, as we have showed that fine-grained blacklisting of attack traffic based on (transport layer) ports is very effective. In turn, detection of legitimate traffic patterns and whitelisting of such patterns during an attack is not possible due to highly variable client traffic.

7.3 RTBH Event Classification

We provide an overview of RTBH event classes in Figure 19. Note that we use the classes introduced in Table 1. In our data set, the major part of RTBH events with DDoS-like anomalies are highly likely to be infrastructure protection RTBHs and represent ≈ 27% of the total events. The potential use of RTBH for prefix squatting protection was found for four ASes and 21 prefixes. The *Other* RTBH events cannot confidently be classified into either use case. We find that for a significant part of these `/32` *other* events, or 13% of the total events, fewer than 10 packets are visible in our data set. Given that some of these prefixes were active through a complete measurement interval, we have to consider that at least a part of these prefixes are not kept intentionally active. Rather, we consider them *RTBH Zombies*, which were once manually triggered but now forgotten. These prefixes pose a risk for their owners, since they are likely to create operational issues for their potential users. For example, connectivity issues of these addresses may be very difficult to debug, since on average, they are only reachable for 50% of the traffic at the IXP.

Finally, 60% of the RTBH events do not match clearly with any common, well-known use case. These events show constant traffic patterns with no anomalous changes. From the classification perspective, this result is not satisfactory and clearly shows the need for further research to completely understand how and why RTBH is used today.

Our results indicate that either not publicly understood use cases of RTBH exist or that the IXP is not a sufficient vantage point to monitor attack traffic. Globally peering ASes might announce RTBHs at all points-of-presences although only a small, local DDoS attack takes place. We emphasize that the presented results are not an artifact of our methodology. Related work shows similar trends with less then 30% RTBHs being related to DDoS attacks. Jonker *et al.,* [14, 16] use a complementary approach to link RTBHs with DDoS attacks. They are utilizing data from an Internet telescope, amplification honeypots, and public BGP route collectors which provide—in contrast to our central vantage point—a distributed view. The authors hypothesize about missed attacks, as their methodology does not allow for the detection of direct and unspoofed attacks. Although being able to observe these additional attack types, we arrive at the same results.

8 CONCLUSIONS AND OUTLOOK

In this paper, we took the first deep dive into the use patterns of remotely triggered blackholing (or BGP blackholing) at a large European Internet Exchange Point. We comprehensively analyzed a data set of control plane correlated with data plane measurements that spanned three months. We did not only consider the behavior of autonomous systems that trigger blackholing but also analyzed networks that received blackhole routes. To our surprise we found several disturbing operational practices which—if improved—could increase the reachability on the Internet infrastructure.

Our further analysis revealed intrinsic measurement challenges for answering important questions about the collateral damage introduced by RTBH. Full packet captures are not available because of privacy and performance reasons, in particular at highly popular IXPs. Therefore, our community relies on packet samples. We found that only a relatively small subset of the captured samples can be used to clearly identify the traffic mix before and during DDoS mitigation, and thus to quantify the collateral damage.

In future work, we will extend our methods to cover a larger portion of RTBH-protected DDoS events when quantifying the collateral damage. We also hope that our results illustrate the potentials and pitfalls of RTBH services to the operator community, which may lead to improved Internet infrastructure security in the mid- to long-term.

Acknowledgments

We would like to thank the anonymous reviewers and our shepherd Cristel Pelsser for their valuable feedback. We also gratefully acknowledge open discussions with the operator community. This work was supported in parts by the German Federal Ministry of Education and Research (BMBF) within the project *X-Check*.

A ETHICAL CONSIDERATIONS

The *control plane* information on remotely triggered blackholes is publicly available at multiple vantage points on the Internet and does not contain potentially privacy-affecting information.

The *sampled flow data* contains data from the network layer and the transport layer. This data potentially contains information that could be correlated or connected to individuals and therefore bears a potential privacy risk. Therefore, the collection and handling of flow data is conducted strictly in accordance to the privacy laws applicable to the collecting organization. The handling of potentially privacy-relevant data is strictly confined to dedicated computer systems that are isolated from the Internet. This data never leaves the premises and control of the collecting organization. Privacy-relevant data is aggregated and anonymized as early as possible in the analysis process. None of the results discussed in this paper can be traced to individual IP addresses or other, privacy-related information.

REFERENCES

[1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. 2012. Anatomy of a Large European IXP. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 163–174.

[2] E. Biersack, Q. Jacquemart, F. Fischer, J. Fuchs, O. Thonnard, G. Theodoridis, D. Tzovaras, and P. Vervier. 2012. Visual analytics for BGP monitoring and prefix hijacking identification. *IEEE Network* 26, 6 (November 2012), 33–39.

[3] J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir. 2014. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 435–448.

[4] Evan Damon, Julian Dale, Evaristo Laron, Jens Mache, Nathan Land, and Richard Weiss. 2012. Hands-on Denial of Service Lab Exercises Using SlowLoris and RUDY. In *Proc. of the Information Security Curriculum Development Conference (InfoSecCD)*. ACM, New York, NY, USA, 21–29.

[5] C. Dietzel, A. Feldmann, and T. King. 2016. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *Proc. of PAM*. Springer Verlag, Berlin, Heidelberg, N.Y., 319–332.

[6] Christoph Dietzel, Matthias Wichtlhuber, Georgios Smaragdakis, and Anja Feldmann. 2018. Stellar: Network Attack Mitigation Using Advanced Blackholing. In *Proc. of ACM CoNEXT*. ACM, New York, NY, USA, 152–164.

[7] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *Proc. of ACM CCS*. ACM, New York, NY, USA, 542–553.

[8] P. Ferguson and D. Senie. 1998. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. RFC 2267. IETF.

[9] David Freedman, Brian Foust, Barry Greene, Ben Maddison, Andrei Robachevsky, Job Snijders, and Sander Steffann. 2018. *Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide*. RIPE Documents ripe-706. RIPE. https://www.ripe.net/publications/docs/ripe-706

[10] V. Giotsas, G. Smaragdakis, C. Dietzel, P. Richter, A. Feldmann, and A. Berger. 2017. Inferring BGP Blackholing Activity in the Internet. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 1–14.

[11] Barry R. Greene. 2019. *Remote Triggered Black Hole (RTBH) Filtering*. Technical Report. http://www.senki.org/operators-security-toolkit/remote-triggered-black-hole-rtbh-filtering/

[12] Nico Hinze, Marcin Nawrocki, Mattijs Jonker, Alberto Dainotti, Thomas C. Schmidt, and Matthias Wählisch. 2018. On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP. In *Proc. of ACM SIGCOMM. Poster Session*. ACM, New York, NY, USA, 57–59.

[13] Patrick Hoffman, Georges Grinstein, and David Pinkney. 1999. Dimensional anchors: A graphic primitive for multidimensional multivariate information visualizations. In *Proc. of ACM NPIV*. ACM, New York, NY, USA, 9–16.

[14] Mattijs Jonker. 2019. A First Joint Look at DoS Attacks and BGP Blackholing in the Wild. RIPE 79 Presentation. https://ripe78.ripe.net/archives/video/22/

[15] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 100–113.

[16] Mattijs Jonker, Aiko Pras, Alberto Dainotti, and Anna Sperotto. 2018. A First Joint Look at DoS Attacks and BGP Blackholing in the Wild. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 457–463.

[17] T. King, C. Dietzel, J. Snijders, G. Doering, and G. Hankins. 2016. *BLACKHOLE Community*. RFC 7999. IETF.

[18] A. Kirkham. 2012. *Issues with Private IP Addressing in the Internet*. RFC 6752. IETF.

[19] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Amplification DDoS Attacks. In *Proc. of RAID*. Springer Verlag, Berlin, Heidelberg, N.Y., 615–636.

[20] B. Krebs. 2016. KrebsOnSecurity Hit With Record DDoS. https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos.

[21] W. Kumari and D. McPherson. 2009. *Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)*. RFC 5635. IETF.

[22] Doug Madory and Tarun Dua. 2014. RE: Prefix hijacking, how to prevent and fix currently. https://seclists.org/nanog/2014/Aug/513

[23] Denis Makrushin. 2017. The cost of launching a DDoS attack. Kaspersky Labs, Online. https://securelist.com/the-cost-of-launching-a-ddos-attack/77784/

[24] P. Marques, N. Sheth, R. Raszuk, B. Greene, J. Mauch, and D. McPherson. 2009. *Dissemination of Flow Specification Rules*. RFC 5575. IETF.

[25] Martin McKeay. 2018. Summer SOTI - DDoS by the numbers. The Akamai Blog. https://blogs.akamai.com/2018/06/summer-soti---ddos-by-the-numbers.html

[26] Loïc Miller and Cristel Pelsser. 2019. A Taxonomy of Attacks using BGP Blackholing. In *Proc. of Computer Security – ESORICS 2019*. Springer Nature, Berlin. LNCS vol. 11735.

[27] Asya Mitseva, Andriy Panchenko, and Thomas Engel. 2018. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications* 124 (2018), 45–60.

[28] C. Morales. 2018. NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us. https://asert.arbornetworks.com/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/.

[29] Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids. 2018. When the Dike Breaks: Dissecting DNS Defenses During DDoS. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 8–21.

[30] C. D. Murta, P. R. Torres Jr., and P. Mohapatra. 2006. QRPp1-4: Characterizing Quality of Time and Topology in a Time Synchronization Network. In *Proc. of IEEE Globecom 2006*. IEEE Press, Piscataway, NJ, USA, 1–5.

[31] Netscout (Arbor). 2019. NETSCOUT Arbor's 13th Annual Worldwide Infrastructure Security Report, Insight into the Global Threat Landscape. Online. https://pages.arbornetworks.com/rs/082-KNA-087/images/13th_Worldwide_Infrastructure_Security_Report.pdf

[32] Alexander Gutnikov Oleg Kupreev, Ekaterina Badovskaya. 2019. DDoS Attacks in Q4 2018. Kaspersky Labs, Online. https://securelist.com/ddos-attacks-in-q4-2018/89565/

[33] Pandas. 2019. Computational Tools. Online. https://pandas.pydata.org/pandas-docs/stable/user_guide/computation.html#exponentially-weighted-windows

[34] Ruoming Pang, Vinod Yegneswaran, Paul Barford, Vern Paxson, and Larry Peterson. 2004. Characteristics of Internet Background Radiation. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 27–40.

[35] PeeringDB. 2019. A freely available, user-maintained, database of networks, and the go-to location for interconnection data. https://www.peeringdb.com/

[36] Dave Piscitello. 2011. IP Prefix Squatting Attacks. https://securityskeptic.typepad.com/the-security-skeptic/2011/06/ip-prefix-squatting-attacks.html

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

[38] Rauchgeist. 2018. DDoS Angriff. Facebook post. https://www.facebook.com/324164061331608/posts/liebe-leute-leider-sind-wir-opfer-einer-ddos-attacke-bitte-habt-keine-sorge-wege/605379806543364/

[39] Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C. Schmidt, and Matthias Wählisch. 2018. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM Sigcomm Computer Communication Review* 48, 1 (January 2018), 19–27.

[40] J. Ryburn. 2015. DDoS Mitigation Using BGP Flowspec. NANOG 63.

[41] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. 2018. ARTEMIS: Neutralizing BGP Hijacking Within a Minute. *IEEE/ACM Trans. Netw.* 26, 6 (Dec. 2018), 2471–2486.

[42] Tomer Shani. 2019. Updated: This DDoS Attack Unleashed the Most Packets Per Second Ever. Here's Why That's Important. Imperva Blog.. https://www.imperva.com/blog/this-ddos-attack-unleashed-the-most-packets-per-second-ever-heres-why-thats-important/

[43] Stephen Strowes. 2019. *Visibility of IPv4 and IPv6 Prefix Lengths in 2019*. Technical Report. RIPE Labs. https://labs.ripe.net/Members/stephen_strowes/visibility-of-prefix-lengths-in-ipv4-and-ipv6

[44] The Verge. 2018. The man behind a spree of gaming network cyberattacks has pleaded guilty. https://www.theverge.com/2018/11/7/18071764/austin-thompson-derptrolling-sony-blizzard-game-ddos-arrest-guilty

[45] Daniel R Thomas, Richard Clayton, and Alastair R Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *Proc. of APWG Symposium on Electronic Crime Research (eCrime)*. IEEE Press, Piscataway, NJ, USA, 79–84.

[46] D. Turk. 2004. *Configuring BGP to Block Denial-of-Service Attacks*. RFC 3882. IETF.

[47] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. 2015. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proc. of NDSS*. ISOC, San Jose, CA, USA, 15.