Peering into the Darkness

The Use of UTRS in Combating DDoS Attacks

Anghel, Radu; Vetrivel, Swaathi; Rodriguez, Elsa Turcios; Sameshima, Kaichi; Makita, Daisuke; Yoshioka, Katsunari; Gañán, Carlos; Zhauniarovich, Yury

DOI
10.1007/978-3-031-51476-0_2

Publication date
2024

Document Version
Final published version

Published in
Computer Security – ESORICS 2023 - 28th European Symposium on Research in Computer Security, The Hague, The Netherlands, September 25–29, 2023, Proceedings

Citation (APA)
Anghel, R., Vetrivel, S., Rodriguez, E. T., Sameshima, K., Makita, D., Yoshioka, K., Gañán, C., & Zhauniarovich, Y. (2024). Peering into the Darkness: The Use of UTRS in Combating DDoS Attacks. In G. Tsudik, M. Conti, K. Liang, & G. Smaragdakis (Eds.), *Computer Security – ESORICS 2023 - 28th European Symposium on Research in Computer Security, The Hague, The Netherlands, September 25–29, 2023, Proceedings* (pp. 23-41). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 14345 LNCS). Springer. https://doi.org/10.1007/978-3-031-51476-0_2

Peering into the Darkness: The Use of UTRS in Combating DDoS Attacks

Radu Anghel[1], Swaathi Vetrivel[1], Elsa Turcios Rodriguez[1], Kaichi Sameshima[2], Daisuke Makita[2,3], Katsunari Yoshioka[2], Carlos Gañán[1], and Yury Zhauniarovich[1(✉)]

[1] Delft University of Technology, Delft, Netherlands
{r.anghel,s.vetrivel,e.r.turciosrodriguez,c.hernandezganan,y.zhauniarovich}@tudelft.nl
[2] Yokohama National University, Yokohama, Japan
{sameshima-kaichi-mx,yoshioka}@ynu.ac.jp
[3] National Institute of Information and Communications Technology, Tokyo, Japan
d.makita@nict.go.jp

Abstract. Remotely Triggered Black Hole (RTBH) is a common DDoS mitigation approach that has been in use for the last two decades. Usually, it is implemented close to the attack victim in networks sharing some type of physical connectivity. The Unwanted Traffic Removal Service (UTRS) project offers a free, global, and relatively low-effort-to-join and operate RTBH alternative by removing the requirement of physical connectivity. Given these unique value propositions of UTRS, this paper aims to understand to what extent UTRS is adopted and used to mitigate DDoS attacks. To reach this goal, we collected two DDoS datasets describing amplification and Internet-of-Things-botnet-driven attacks and correlated them with the information from the third dataset containing blackholing requests propagated to the members of UTRS. Our findings suggest that, currently, just a small portion of UTRS members (approximately 10%) trigger mitigation attempts: out of 1200+ UTRS members, only 124 triggered blackholing events during our study. Among those, with high probability, 25 Autonomous Systems (ASes) reacted on AmpPot attacks mitigating 0.025% of them globally or 1.03% targeting UTRS members; 2 countered IoT-botnet-driven attacks alleviating 0.001% of them globally or 0.06% targeting UTRS members. This suggests that UTRS can be a useful tool in mitigating DDoS attacks, but it is not widely used.

Keywords: UTRS · RTBH · DDoS attacks

1 Introduction

Distributed Denial-of-Service (DDoS) attacks are on the rise [10], and they are proving difficult to defend against. DDoS attacks aim to bring down a targeted server, service, or network by overwhelming the target system or its network infrastructure with traffic [22]. DDoS attacks have financial costs to victims,

G. Tsudik et al. (Eds.): ESORICS 2023, LNCS 14345, pp. 23–41, 2024.
https://doi.org/10.1007/978-3-031-51476-0_2

as they could lose thousands of dollars and possibly go out of business [10]. Attackers employ different techniques and tools to launch them. Lately, massive DDoS attacks are often fueled by embedded and IoT devices [4] and widely leverage reflectors to amplify the volume of traffic directed at victims [2].

To mitigate DDoS attacks, network operators and service providers have at their disposal various countermeasures and mitigation strategies. Among some of the most common countermeasures are clean pipe, content delivery network (CDN) attack dilution, antiDDoS proxy, and Remotely Triggered Black Hole (RTBH) service [9,14,20]. The latter has become popular and used by Internet Service Providers (ISPs) for about two decades [25]. RTBH leverages the Border Gateway Protocol (BGP) to reroute attack traffic to places that minimize harm, typically by dropping it [25].

There are two approaches to putting RTBH into practice. The first one relies on the direct connection between Autonomous Systems with customer-provider or peering relationships. The second approach employs an Internet Exchange Point (IXP), e.g., DE-CIX [8] or Equinix [11], that acts as a central point distributing blackholing requests to their participants. Regardless of the chosen approach, RTBH requires physical connectivity from all participants. Moreover, IXP-provided RTHB is typically a paid service [3].

For those network operators that do not have access to IXPs offering RTBH or the service is too expensive, there is at least one free alternative, namely joining an RTBH project. An example of such a project is the Unwanted Traffic Removal Service (UTRS) [17], a community project of Team Cymru. To the best of our knowledge, UTRS is currently the *only global free-to-participate RTBH initiative*. UTRS is essentially an RTBH operated by a trusted third party, similar to RTBHs operated by IXPs. In this case, instead of an IXP, UTRS acts as the central point that decides whether to accept and distribute the blackholing requests to its participants. Because it does not need the physical connectivity of participants, UTRS is not limited geographically.

UTRS is an initiative that any network can join with relatively low effort. As stated on the project webpage [23], UTRS is offered as a free service to owners of unique Autonomous System Numbers (ASNs). Joining the UTRS project is done by filling out an online form that is manually checked by Team Cymru to validate ownership/access of the AS. If the validation is successful, Team Cymru provides BGP configuration details and generic guides for various router types [24]. As of August 2021, 1200+ networks were participating in this initiative [7].

In contrast to an RTBH, UTRS helps to mitigate DDoS attacks *closer to their source* [23]. When implementing RTBH over direct connections such as with upstream ISPs, peering relationships or IXPs, the blackholing happens close to the attack destination, still causing congestion in networks transiting the attack. With UTRS, however, the participants could be all over the world, that, in theory, improves the chance of stopping an attack closer to its source. However, the likelihood that DDoS attacks are stopped near their source is diminished if the number of participants is low. Thus, it is important to incentivize ASes to join such initiatives to decrease transition junk traffic [26].

In this work, we focus on understanding to what extent UTRS is used globally among the ASes to mitigate attacks and which DDoS attacks trigger mitigation attempts via UTRS. To the best of our knowledge, we are *the first who study adoption and the effectiveness of this service in detail*. The closest work [13] has analyzed DDoS attacks and compared them to BGP blackholing events. The information provided in our work may be valuable for network operators who would like to asses joining UTRS as an alternative to the RTBHs offered through direct connectivity.

To reach our goal, we collect six months of blackholed prefixes from UTRS at five minutes intervals and compare them to two different data sources that provide information about DDoS targets. The first source is an IoT botnets Milker service designed to extract information about targets of IoT-driven DDoS attacks. It gathers IoT botnet malware samples, analyzes them, and then behaves as one of the bots, enabling us to gather DDoS attack commands from C&C servers. The second, AmpPot [16], tracks reflection and amplification attacks. We analyze the data collected from these sources to identify trends and patterns in using UTRS to mitigate DDoS attacks. We set up ourselves to answer the following research questions:

(**RQ1**): How many UTRS members use this service to mitigate attacks?
(**RQ2**): To what extent are DDoS attacks triggering mitigation attempts via UTRS?
(**RQ3**): To what extent can UTRS announcements be explained by amplification DDoS attacks?
(**RQ4**): To what extent can UTRS announcements be explained by IoT-botnet-driven DDoS attacks?

The contributions of this paper are the following:

- Our analysis provides valuable insights into the adoption of the UTRS service. It shows that only around 10% of all UTRS members actively use this service to mitigate attacks.
- We provide a comprehensive analysis of the usage of UTRS in mitigating DDoS attacks coming from IoT botnets and amplification attacks. Our results show that while UTRS is used to mitigate reflection attacks, it is barely triggered to handle IoT-botnet DDoS attacks.
- We identify trends and patterns in the use of UTRS to mitigate DDoS attacks.

2 Related Work

To combat DDoS attacks, various mitigation methods have been suggested, such as blackholing and rate limiting. In this section, we examine the current research on blackholing, a technique that discards traffic directed at a targeted IP address, preventing it from reaching the victim. We focus on recent studies investigating the effectiveness of blackholing, its limitations, and proposed enhancements to

mitigate its drawbacks. We also discuss related work on detecting and analyzing blackholing events and their correlation with DDoS attacks.

In the closest work, Jonker et al. [13] correlate BGP blackholing events with the data from two attack datasets: "Randomly and Uniformly Spoofed Attacks" collected from CAIDA's telescope and AmpPot. The blackholing events are inferred by applying some heuristics to public BGP data collected by RouteViews and RIPE RIS. An attack is considered to be blackholed only if its target IP address is within the BGP announced network and happens no more than 24 h prior to the announcement. Their results show that only 456k of the 28.16M attacks (1.62%) from CAIDA/AmpPot datasets are blackholed, involving only 0.81% of all uniquely targeted IP addresses. Another finding shows that for attacks found in both CAIDA and AmpPot (447.6k attacks), 18.4k (4.12%) involving 5.7k (3.25%) unique IPs are blackholed, suggesting that more serious attacks are more likely to be blackholed.

In [26], the authors evaluated how much junk traffic is generated and transferred by ISPs due to amplification attacks. They proposed a method to filter this traffic out by using an AmpPot-like honeypot to obtain information about the victims and Software Defined Network (SDN) to block the traffic targeting those victims. Similarly, UTRS can be used instead of SDN.

Giotsas et al. [12] developed a methodology to detect BGP blackholing activity using datasets from RIPE RIS, RouteViews, PCH, and a "CDN." The findings reveal that 26 IXPs and 242 networks offer RTBH services to customers, peers, and members, and 96.64% of RTBH events are for IPv4, for the measurement period between August 2016 - March 2017.

Nawrocki et al. [21] analyze RTBH events at a large European IXP, and correlate the BGP data with IPFIX flows, which capture information about actual traffic passing through the IXP. The paper finds that more than 95% of traffic and more than 98% of RTBH events are for IPv4 addresses. The IXP had 830 member ASes on average during the measurement period (104 days), of which 78 member ASes announced 1107 IPs to be blackholed for 170 origin ASes (some customers of the 78 members). The events were repeated, indicating that the same IP was under different attacks at different times.

In [15], DDoS attack types at a major IXP from September 2019 to April 2020 are analyzed. The authors find that 89.9% of the attacks use known amplification protocols such as DNS and NTP and IP spoofing. The data is compared to a "commercial world-wide honeypot network," revealing a correlation with only 8% of the attacks observed at the IXP, covering only 33% of target IPs.

Finally, the authors in [9] introduced the "Advanced Blackholing" system that was tested at a major IXP. The system provides improved granularity and the authors claimed that it can successfully mitigate attacks without disrupting the service to the victim. The "Stellar Advanced Blackholing" combines the benefits of RTBH, FlowSpec, ACL filters, and Traffic Scrubbing, without requiring cooperation between networks participating in the IXP. The signaling in Stellar uses BGP, similar to RTBH, but the filters are deployed at the IXP level by utilizing OpenFlow, thus relying on the IXP for both the software and hardware implementation.

Table 1. Datasets description

Dataset	# entries	# targets	# unique target IPs	Duration (sec)		
				min	mean	max
UTRS	533,257	7,820	7,830	300.0	4,682.7	413,700.0
AmpPot	1,616,184	1,080,770	1,080,770	0.5	891.5	1,949,571.0
Milker	223,267	46,764	2,787,522	1.0	93.0	3,600.0

3 Data Collection and Datasets Descriptives

Our study aims to determine to what extent DDoS attacks lead to UTRS participants triggering mitigation attempts. To achieve this goal, from October 2022 to April 2023 (6 months), we gathered three datasets: UTRS, IoT Milker, and AmpPot. Table 1 describes the collected datasets. The number of entries per dataset presents the total number of observations per day in each data set aggregated for the whole period of observation. Next, the column number of targets presents the number of target networks in the case of UTRS and Milker, while for AmpPot the number displayed is the number of individual IPs targeted. The column number of unique targeted IPs is self-explanatory. Finally, the table shows the minimum (min), mean, and maximum (max) duration of attacks.

By combining these datasets, we aim to provide a comprehensive analysis of how amplification and IoT-botnet DDoS attacks influence UTRS participants to trigger UTRS mitigation attempts. In this section, we describe the methodology used to collect and analyze each of these datasets.

3.1 UTRS

To collect the UTRS data set, we registered our own AS and joined the UTRS project. This allowed us to receive the BGP announcements to black hole traffic to particular targets[1] from this service. We made snapshots of active BGP routes every 5 min. Then, we stitched the data using the following rule: if the same target appears in several consecutive snapshots, we assume that it was blocked during the whole period of time[2]. As a result, we get a dataset that consists of targets with start and end blackholing times with 5 min granularity.

Our UTRS dataset has 533, 257 entries. Figure 1a shows entries distribution per day (hereafter, the UTRS dataset lines are colored orange). On average, we get 3, 122 entries each day in the UTRS dataset, with the maximum reaching 9, 427 entries.

In the dataset, there are 7, 820 unique target networks (see Table 1) from 124 ASes that contain 7, 830 individual IP addresses. The majority of UTRS

[1] UTRS members can announce up to a /25 of IPv4 addresses and up to /49 for IPv6 from their ASes as targets.

[2] If a target network is added and removed within a 5-minute interval, it will not appear in any dump and we will not record it.

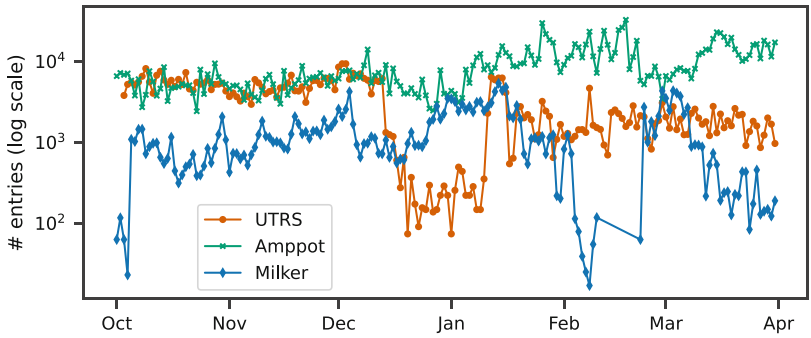

(a) Number of entries in the datasets per day

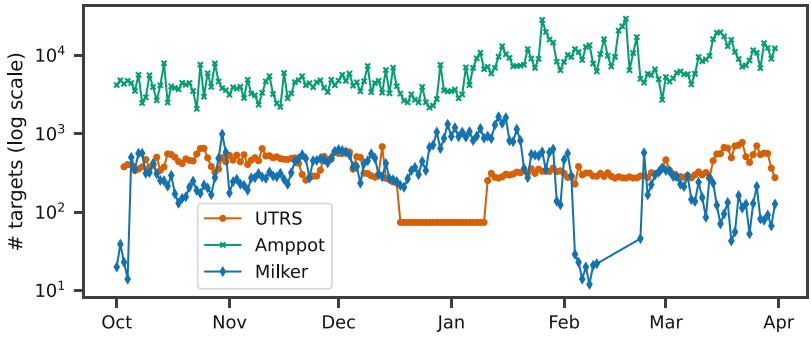

(b) Number of targets in the datasets per day

Fig. 1. Daily datasets characteristics

announcements target individual IP addresses (/32 prefix length). Only 2 entries within the 6-month observation period target the same /27 subnetwork on the same day. The number of targets in the UTRS dataset varies greatly from day to day (see Fig. 1b). The maximum number of observed IPs during one day was 776, and the minimum of 74. The mean is 357 IPs per day, with a standard deviation of 159.9. In terms of trends, there is no clear pattern in the data. At the same time, it is obvious that the number of targets in the months of 2022 seems to be consistently higher than in the ones of 2023.

There is a cluster of days between December 18th and January 11th where the number of targets is relatively low (74 unique targets). Interestingly, during the same period, the number of entries in the dataset varies from 74 to 654 (see Fig. 1a), but the same 74 targets are announced during these days. We suspect the service was not functioning properly because our test announcements were not propagated through the service during that period.

Figure 2 shows the empirical Cumulative Distribution Function (eCDF) and the corresponding Kernel Density Estimation (KDE) of durations. As we can see

from this figure, the largest part of the UTRS announcements (around 21%) has a duration equal to 5 min, which is connected with the peculiarities of the UTRS dataset collection. This tells us that a huge portion of the UTRS announcements lasts less than 5-min, and we might miss them dumping BGP tables using this period. The longest duration of announcements in the UTRS dataset was 4 days, 18 h and 55 min (or 413, 700 s) happened at the end of October 2022, with the mean and the median of 4, 682.7 and 1, 798.0 s correspondingly.

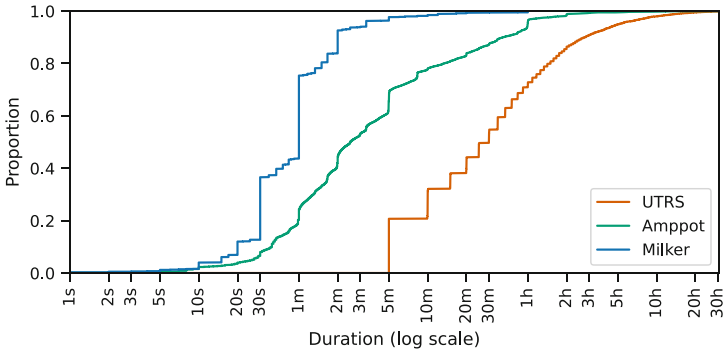

(a) eCDF

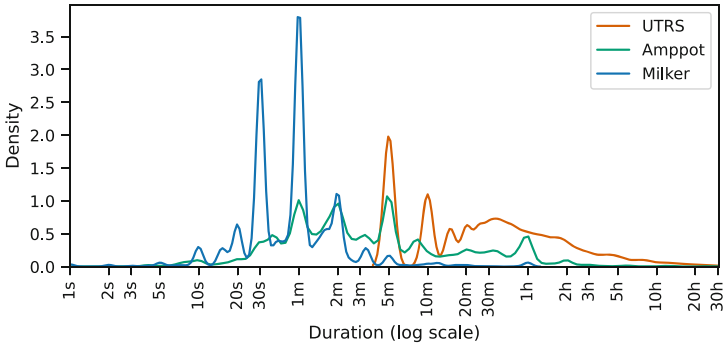

(b) KDE

Fig. 2. Events duration characteristics (axis x is in logarithmic scale)

3.2 AmpPot

As the second dataset, we gathered attack data provided by the AmpPot project[3], a reflection and amplification honeypot [16]. The AmpPot honeypot tracks reflection and amplification attacks by mimicking vulnerable protocols, such as DNS, NTP, and SSDP. This allows the honeypot to be discovered by

[3] https://sec.ynu.codes/dos.

attackers scanning for reflectors and subsequently used in their attacks. By logging the received requests, the AmpPot data allows for the inference of various information about an attack, including the victim's IP address and destination port, the start and end time of the attack, and the attack intensity.

We collected the amplification attack data from 19 AmpPot sensors (11 proxied and 8 agnostic, see [16] for details). As it was previously shown, e.g., in [18], AmpPot attacks recorded by different sensors are often related. There are several strategies for how to combine attacks recorded by different AmpPot sensors. In this work, we use the same approach as in [16]: we stitch attacks seen by multiple sensors into one combined attack if the target and the destination port are the same and there is less than an hour time difference between the attacks. As a result, we got a dataset that contains 1,616,184 attack entries (see Table 1).

The number of entries ranges from 2,435 to 32,908, with a mean of 8,969.7 attacks per day. During this period, 1,080,770 unique IPs[4] from 15,825 ASes were attacked. The count of unique IPs per day varies from 2,082 to 29,402, with a mean of 6,943.6. Figure 1a shows the count of unique attack entries per day over the entire observation period, while Fig. 1b reports the number of targets (hereafter, the AmpPot dataset lines are painted green).

The mean attack duration is 891.5 s, and the median is 150.0 s. Figure 2 shows the eCDF and KDE plots of the duration of AmpPot attacks. As we can see on the KDE plot, there are clear peaks at 1, 2 and 5 min, which show the typical durations of AmpPot attacks. The longest recorded AmpPot attack lasts for 22 days 13 h 32 min and 51 s (1,949,571 s), which is probably a routine scan.

3.3 IoT Milker

We set up a C&C monitoring system to monitor IoT DDoS attacks. Our C&C monitoring system consists of three main components: the sandbox, the C&C identifier, and the C&C milker. The sandbox safely executes IoT malware and observes its communications with C&C servers.

The C&C identifier part detects and identifies C&C servers using a two-step process. In the first step, we use frequency analysis to look for frequent access to certain IP addresses and ports, which can indicate the presence of a C&C server. This method does not rely on signatures or other forms of static analysis, making it less susceptible to evasion by IoT botnets. The second step uses predefined rules, manually prepared by analysts, to identify the C&C protocol based on the first payload sent by the sandboxed malware. The initial set of rules is based on the Mirai source code, but new rules are added as needed. Each rule is associated with a corresponding milker script, which imitates the identified C&C protocol and allows the observation system to capture the real-time activities of the IoT botnet.

The Milker script is a manually created script that imitates the IoT bot behavior. Using the identified protocol, it connects to the combinations of IP

[4] For this dataset, the number of targets corresponds to the number of IPs because AmpPot records attacks to individual IPs rather networks.

addresses and ports of C&C servers detected with the frequency analysis and monitors the commands sent by the C&C servers. Thus, using the milker script, we are able to capture the real-time activities of IoT botnets, including the commands to DDoS particular victims. For this work, we collect the start time of a DDoS attack, the target IP and port, and its duration. Since we collect the commands sent by the C&C, we do not typically have other statistics on the attack, like the total number of packets and packets per second.

Over the entire collection period, we gathered a dataset containing 223,270 entries. The majority of entries (217,565) attack individual IPs (have /32 prefix length). At the same time, there are several cases that attack large subnetworks: 1 attack targets /2, and 2 – /4. Such distributed attacks do not make sense; thus, we assumed that botnet owners initiated them due to mistyping. Most likely, in one /2 and one /4 cases, the prefix length was intended as /24 by the attacker. We decided to remove all these 3 entries from the dataset in order to not overcount potential unique IP addresses, that would also overlap with loopback ranges. As a result, we get a dataset that contains 223,267 entries, targeting 46,764 unique victim networks with 2,787,522 unique IP addresses.

Figure 1a and Fig. 1b show the number of entries and targets per day (hereafter, the Milker dataset line has a blue color). On average, we observed 1,306.7 entries per day targeting 399.2 networks, with the highest values being 5,396 attacks and 1,648 unique targets and the lowest 17 entries and 12 networks per day correspondingly. Note that we do not have Milker data for 11^{th}–22^{nd} February 2023, due to our infrastructure upgrade.

The duration of the Milker attacks, as observed by the value sent by the C&C to the bots, ranges from 1 second to a maximum of one hour (see Fig. 2). The upper limit of 3,600 s is probably an artifact of the Mirai code[5] which raises an error if the duration is longer than an hour. As we can see in the figure, the majority of attacks in the Milker dataset have 30 s, 1 or 2 min durations.

4 Findings

The two collected DDoS datasets, AmpPot and Milker, are the sources of information about amplification and IoT-botnet DDoS attacks correspondingly. For each DDoS dataset, we compute two views representing the intersections with the UTRS dataset data. The first view (*Exact Interval (EI)*) contains data about exact time interval intersections between DDoS attacks and UTRS announcements, while the second (*Offset Interval (OI)*) represents the intersections of DDoS attack intervals with the UTRS announcement span extended by 12 h in both directions. The reason for computing the view with Offset Interval is the following. A sensor can record an attack that may be already over before it is blocked by UTRS. Similarly, some sensors can still continue registering attack packets even though the corresponding entry is removed from the UTRS table.

[5] https://github.com/jgamblin/Mirai-Source-Code/blob/master/mirai/cnc/attack.go.

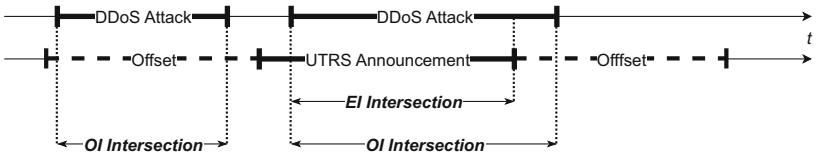

Fig. 3. Intersections views (EI - Exact Interval, OI - Offset Interval)

As we described in Sect. 3, both datasets, AmpPot and Milker, contain information about the victims (individual IPs in the case of AmpPot and prefixes for the Milker dataset). To build the views described above, we first compute a map relating DDoS attack victims with the UTRS announcement targets. Then, for each UTRS target, we search for all intersections of UTRS announcement spans with the time intervals of DDoS attacks targeting the same victim. In the case of *Exact Interval (EI)*, we look for exact intersections of the time intervals. For *Offset Interval (OI)*, we, at first, expand the span of the UTRS announcement in both directions by 12 h and then find the intersections of the attack's time intervals with this expanded period. Figure 3 shows the relative position of these intervals graphically and explains the concepts visually. As a result, we get four views: two (EI and OI) for the UTRS-AmpPot datasets and another two for the UTRS-Milker datasets.

4.1 General Overview

Table 2 reports the characteristics of the computed views. The first row, "# of entries", contains the number of entries for the corresponding computed view. As expected, the OI views contain considerably more entries. Interestingly, the number of entries in the UTRS-Milker view is significantly lower than in UTRS-AmpPot, although the number of target IPs in the former is higher than in the latter (see Table 1). The reason is that the number of unique targets is lower in the UTRS-Milker view. Additionally, the AmpPot attacks are longer on average, which increases the probability of the intersection with a UTRS announcement time interval. The row "# of unique UTRS targets" reports the number of unique targets in the UTRS announcements. For UTRS-AmpPot, the values in this row are equal to the ones in the "# of unique DDoS attack targets" row, meaning that one AmpPot attack is typically covered by only one UTRS announcement. At the same time, for UTRS-Milker, a Milker attack may trigger several UTRS announcements because some Milker attacks target networks rather than individual IPs. The values in the last row, "Mean entries # per UTRS announcement", show that one UTRS announcement covers more than one attack. For instance, every UTRS announcement has an EI intersection with about 1.55 AmpPot and 1.12 Milker attacks correspondingly.

Table 2. Views characteristics (EI - Exact Interval, OI - Offset Interval)

Parameter	UTRS-AmpPot		UTRS-Milker	
	EI	OI	EI	OI
# of entries	468	6,774	9	791
# of unique DDoS attack targets	249	1,268	2	143
# of unique UTRS targets	249	1,268	8	163
# of unique UTRS ASNs	25	43	2	6
Mean entries # per UTRS announcement	1.55	1.76	1.12	1.88

The "# of unique UTRS ASNs" row shows the number of ASes that launched the UTRS announcements in the corresponding views. With high probability, we can assume that only 25 and 2 ASes reacted to AmpPot and Milker attacks correspondingly by commencing UTRS announcements (EI intersections), while 43 and 6 of ASes are triggered by AmpPot and Milker attacks with lower probability (OI intersections).

Among those, there are 43 unique ASes. Table 3 (in appendix) lists the ASes (in the anonymized form) and what countries they belong to[6]. More than a quarter (11 out of 43) of all the ASes are from Brazil. While this may seem unexpected, a recent study [19] shows that Brazilian ISPs adopt anti-DDoS security best practices (e.g., source address validation) significantly faster than the providers in the rest of the world. Thus, it is highly likely that they have also employed the protection provided by UTRS. Other prominent countries where operators use UTRS to protect against DDoS attacks are the USA (21% or 9 ASes) and Argentina (16% or 7 instances).

Table 3 also reports the total number of AmpPot and Milker attacks and how many of those are mitigated with the help of UTRS per each AS individually. On average, only 1.03% of AmpPot and 0.06% of Milker attacks on the UTRS members trigger the announcements for EI, while for OI, those figures are equal to 8.86% and 6.88% correspondingly.

As for the absolute numbers, only 0.025% and 0.212% of all AmpPot attacks trigger UTRS announcements for EI and OI correspondingly. These numbers are considerably lower than the ones reported for AmpPot dataset in the work by Jonker et al. [13]. According to their measurements, 1.97% of all AmpPot attacks triggered BGP blackholing. There are two reasons for this difference. First, UTRS-based blackholing events represent only a tiny subset of all BGP blackholing events. Second, the authors use 24 h interval prior to a blackholing event to match the attacks, while we consider only 12 h. Moreover, the authors relied on the assumption that all BGP announcements targeting networks smaller or equal to /24 correspond to BGP blackholing events, that can be an optimistic overapproximation. In our work, we do not need to make such an assumption because we know exactly what BGP announcements are trig-

[6] The ASes' country codes are obtained using the Caida's AS Rank [6] dataset.

gered by UTRS. As for the Milker attacks, the numbers are even lower: only 0.001% and 0.147% of all attacks trigger UTRS announcements for EI and OI correspondingly.

4.2 Time Lags

Given the obtained views, we can analyze time lags between UTRS and DDoS attack events. Based on the timestamps in the UTRS-AmpPot and UTRS-Milker EI and OI views, we calculate time difference between the UTRS and the corresponding DDoS attack start and end events. Figure 4 shows the eCDFs of these time lags: a) for Exact Interval views, b) for Offset Interval views. The positive lag (values on the x axis) means that the UTRS event has happened later than the corresponding DDoS attack event.

For EI, as we can see in Fig. 4a, around 67% of all AmpPot attacks start before the UTRS announcement, which is expected. At the same time, around 32% of AmpPot attacks start after the corresponding UTRS announcement. One potential explanation is that attackers use different amplifiers, and our AmpPot sensors are exploited at later stages. Moreover, due to the peculiarities of our UTRS dataset collection, the reported time of a UTRS announcement is shifted into the future compared to the actual time. As expected, for a large portion of AmpPot attacks (around 79%), the corresponding target is removed from the UTRS BGP table after the attack is over. However, the peculiarities of the UTRS dataset collection may have a negative effect on such a high percentage. For 100% of all entries in the UTRS-Milker EI view, announcements happen $1-3$ h before the start of the corresponding attacks recorded by Milker. However, the size of the UTRS-Milker view is very small (only 9 entries), so the results may be nonrepresentative.

Considering OI views (see Fig. 4b), as expected, the majority of UTRS announcements (around 83%) and removals (roughly 83%) happen after the AmpPot attack start and end events correspondingly. At the same time, the same figures for the UTRS-Milker data constitute only 13% and 14% correspondingly. That indicates Milker attacks are highly unlikely caused by the corresponding UTRS announcement.

4.3 Characterization of Blackholed Attacks

To study the characteristics of the attacks that drive the UTRS requests, we examine the properties, namely duration, the total number of packets, and their intensity. Note that we can do this study only for the AmpPot attacks because the amplification honeypots collect the attack metrics, such as the number of packets and the duration of the attacks. Unfortunately, the Milker monitor does not have access to such metrics[7]. Additionally, we consider only the Exact Interval

[7] The Milker monitor registers attack durations, but those values are fixed by the IoT malware owners.

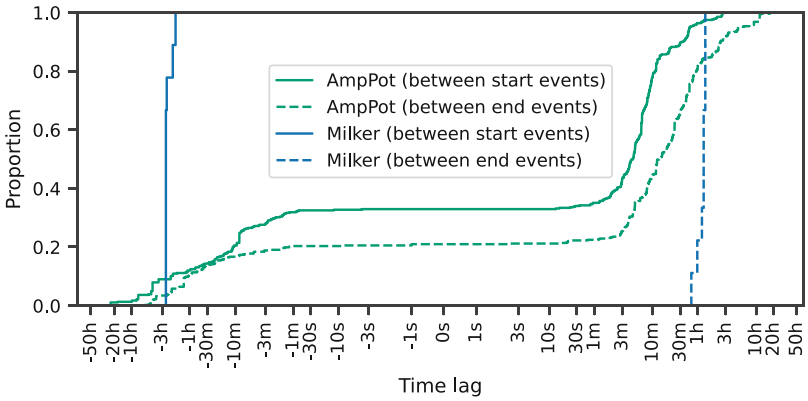

(a) Exact Interval

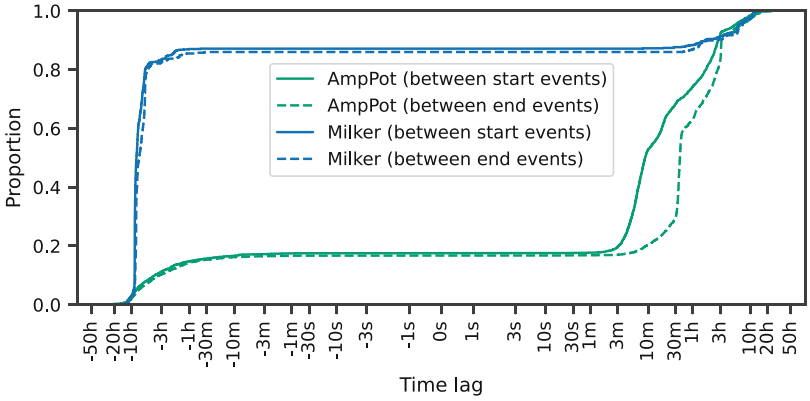

(b) Offset Interval

Fig. 4. Time lags between UTRS events and corresponding DDoS attack events (x axis has a symmetrical log scale)

intersections because the UTRS requests are more likely to be driven by the corresponding AmpPot attacks.

In this section, we compare the properties of the attacks that target only the ASes that trigger UTRS announcements (mitigated/blackholed by UTRS). To build the dataset for such analysis, we extracted all the attacks from the AmpPot dataset (see Sect. 3.2) that target 25 ASNs from the UTRS-AmpPot view (see Table 2). Then, we mark the attacks found in the UTRS-AmpPot view as *blackholed*. Thus, our dataset contains 38,251 entries overall, out of which 398 are marked as blackholed. Note that our dataset is highly unbalanced: only around 1% of all entries belong to the blackholed class, while the rest 99% are members of the non-blackholed class. Luckily, eCDFs, which we use in

this section, accurately reflect the underlying distribution of the observed data, regardless of whether the dataset is balanced or not.

Figure 5a shows the eCDFs of the AmpPot attack duration. Hereafter, the solid line corresponds to all attack data (*Overall*), while the dashed one reflects the information of the attacks marked as blackholed (*Blackholed*). As we can see, the dashed line lies under the solid, showing that the attacks triggering UTRS generally last longer. This result is easy to explain because ISPs prefer to activate UTRS only for long-lasting attacks rather than short-lived ones. Our results also coincide with the ones reported by Jonker et al. [13], who also found that the AmpPot attacks mitigated using BGP blackholing last longer.

Figure 5b shows the eCDFs of the total number of packets recorded by AmpPot sensors. As we can see, the dashed line lies under the solid, meaning that UTRS-mitigated attacks have a higher volume. This result is also expected — indeed, it is more likely for an ISP to mitigate more volumetric attacks.

Figure 5c shows the eCDFs of the intensity of the attacks — mean number of packets per second (pps). Note that this parameter is obtained as the total number of packets divided by the duration of the corresponding stitched attack. Thus, this value correlates with the intensity of the attacks on a victim. As we can see, below the speed of 10 pps, the dashed line lies under the solid one, showing that the blackholed attacks are more intense below this speed. However, after 10 pps, the lines switch. While it is harder to explain, in our opinion, such behavior may be due to the victim ISPs (and their allies) starting to activate multiple protection mechanisms, some of which may also include filtering of attack packets to amplifiers.

5 Limitations and Future Work

In this study, we investigated the use of UTRS as a mechanism to mitigate DDoS attacks. However, there are certain limitations to our findings that need to be acknowledged.

Firstly, our study only captures data on the use of UTRS, and we do not have information on whether network operators used other mechanisms in conjunction with UTRS to mitigate DDoS attacks, or if the attack was successfully mitigated. This means that we cannot accurately quantify the effectiveness of UTRS as a standalone mechanism.

Secondly, our attack data only includes information from AmpPot and IoT Milker. Therefore, we may have missed attacks that were not detected by these sources. To address this limitation, we suggest future work to investigate the effectiveness of UTRS in conjunction with other DDoS attack data sources.

Thirdly, we collected the UTRS data as periodic snapshots every 5 min. We collected the data this way instead of logging all BGP updates due to the constant errors in the BGP sessions with the UTRS project that caused the recorded events to be mostly `announces` with few `withdraws`, thus making it impossible to track when a blackholing event stops. However, if an announcement happens within this interval, we will not be able to record it. This means we may miss

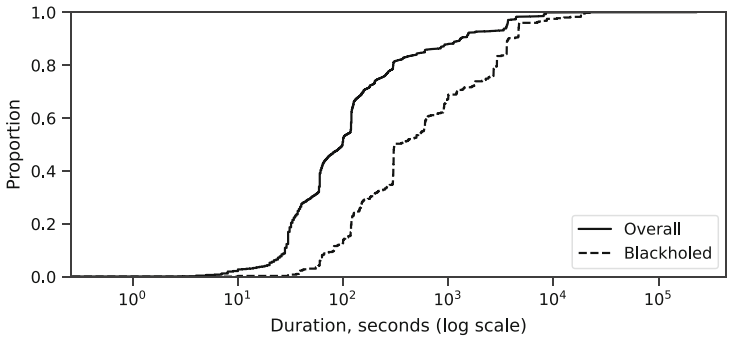

(a) Attack duration

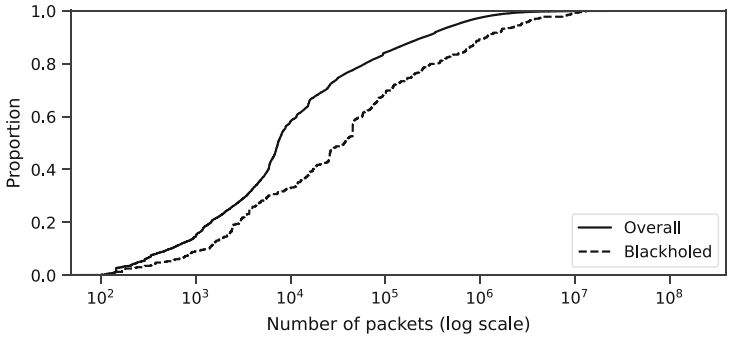

(b) Total number of packets

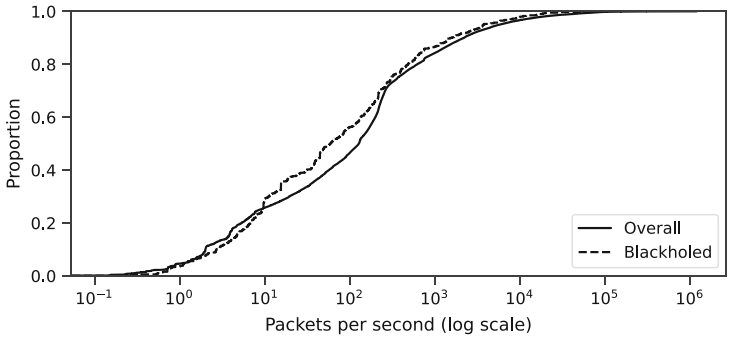

(c) Mean number of packets per second

Fig. 5. Blackholed and non-blackholed attack characteristics eCDFs

some of the announcements, thus potentially underestimating the number of mitigated DDoS attacks and ASes making the announcements. Currently, we are exploring the feasibility of collecting data about the UTRS announcements as soon as they are propagated to the service members.

Lastly, we cannot confirm whether UTRS participants actually blocked the malicious traffic or not. We can only infer that UTRS was used based on the announcement of blocked IP addresses.

6 Ethics

The research presented in this paper was conducted in accordance with the ethical standards set forth by the Institutional Review Board (IRB) of our institution. The data used in this study was obtained through active and passive internet measurements and monitoring botnet commands, which were carried out in compliance with the principles outlined in the Menlo Report [5]. We took steps to ensure that our monitoring did not interfere with the normal functioning of any network and did not compromise the privacy or security of any operators.

7 Conclusions

DDoS attacks continue to remain a significant threat to the Internet. RTBH is one of the most popular and effective mechanisms to mitigate this threat used by ISPs for about two decades [25]. Unfortunately, it requires physical connectivity between participants and may incur additional costs if provided by an IXP. UTRS is a free, global, and low-effort-to-join alternative to RTBH. Given these unique advantages, the goal of this work is to shed more light on this service. In particular, we aim to investigate how many UTRS members use this service to mitigate amplification and IoT-botnet-driven DDoS attacks.

The results of our analysis show that only 0.025% of amplification and 0.001% of IoT-botnet-driven attacks are highly likely mitigated using UTRS. Among UTRS members, only 124 ASes (around 10%) actively use this service to mitigate attacks. Such low adoption means that it might not be that effective. Indeed, currently, only about 1% of all assigned ASNs [1] are UTRS members [7], and an even smaller percentage abides the blackholing requests sent through this service. Thus, we should incentivize ASes to join UTRS and other similar initiatives to increase the effectiveness of this technology and to protect the Internet from DDoS attacks.

Acknowledgements. This work is partly supported by the Dutch Research Council (NWO) under the RAPID project (Grant No. CS.007), by the MITIGATE project (JPJ000254) supported by MIC, Japan, and the commissioned research (No.05201) by NICT. This work was also supported by JSPS KAKENHI Grant Numbers 21H03444 and 21KK0178.

Appendix A: List of ASes Mitigating DDoS Attacks

Table 3. ASes mitigating DDoS attacks

AS	CC	AmpPot					Milker				
		Total attacks	EI mitigated		OI mitigated		Total attacks	EI mitigated		OI mitigated	
			#	%	#	%		#	%	#	%
UAO	AR	2	2	100.00	2	100.00	0	0	0.00	0	0.00
GNX	ES	1	1	100.00	1	100.00	0	0	0.00	0	0.00
VYC	DE	6	5	83.33	6	100.00	0	0	0.00	0	0.00
SMI	BY	5	3	60.00	4	80.00	0	0	0.00	0	0.00
H2W	AR	4	2	50.00	4	100.00	0	0	0.00	0	0.00
BME	AR	135	45	33.33	78	57.78	0	0	0.00	0	0.00
WO6	FR	99	30	30.30	30	30.30	0	0	0.00	0	0.00
WYU	AR	4	1	25.00	3	75.00	0	0	0.00	0	0.00
FAX	AR	15	3	20.00	3	20.00	0	0	0.00	0	0.00
MV3	TR	25	5	20.00	5	20.00	0	0	0.00	0	0.00
KIM	KH	15	2	13.33	6	40.00	0	0	0.00	0	0.00
44I	CA	24	2	8.33	7	29.17	1	0	0.00	0	0.00
FCF	BR	139	9	6.47	55	39.57	6	0	0.00	0	0.00
L7L	PL	173	9	5.20	66	38.15	38	0	0.00	1	2.63
IVB	AU	98	5	5.10	29	29.59	12	0	0.00	8	66.67
SHB	BR	315	11	3.49	20	6.35	0	0	0.00	0	0.00
QAM	US	124	4	3.23	33	26.61	46	0	0.00	22	47.83
73H	IE	38	1	2.63	7	18.42	3	0	0.00	0	0.00
564	US	6,140	**141**	2.30	743	12.10	819	**2**	0.24	116	14.16
QWW	BR	1,252	22	1.76	22	1.76	0	0	0.00	0	0.00
PZS	US	570	8	1.40	22	3.86	14	0	0.00	1	7.14
T7M	BR	1,391	8	0.58	9	0.65	0	0	0.00	0	0.00
JA4	US	**21,615**	77	0.36	**2,229**	10.31	**3,835**	1	0.03	**182**	4.75
ZNH	BR	308	1	0.32	1	0.32	0	0	0.00	0	0.00
RMF	IR	5,753	1	0.02	1	0.02	22	0	0.00	0	0.00
Z76	BR	2	0	0.00	2	100.00	0	0	0.00	0	0.00
4T2	PL	53	0	0.00	1	1.89	4	0	0.00	0	0.00
OVZ	US	2	0	0.00	1	50.00	0	0	0.00	0	0.00
RWB	GB	14	0	0.00	7	50.00	0	0	0.00	0	0.00
G6P	US	10	0	0.00	1	10.00	0	0	0.00	0	0.00
PXB	AR	32	0	0.00	1	3.12	0	0	0.00	0	0.00
63A	BR	7	0	0.00	4	57.14	0	0	0.00	0	0.00
6KY	AT	2	0	0.00	2	100.00	0	0	0.00	0	0.00
JUX	BR	3	0	0.00	1	33.33	0	0	0.00	0	0.00
XMN	AR	3	0	0.00	3	100.00	0	0	0.00	0	0.00
RQI	US	2	0	0.00	1	50.00	0	0	0.00	0	0.00
L4K	SG	3	0	0.00	2	66.67	0	0	0.00	0	0.00
3L4	PK	7	0	0.00	1	14.29	0	0	0.00	0	0.00
VJC	US	2	0	0.00	2	100.00	0	0	0.00	0	0.00
U75	BR	18	0	0.00	3	16.67	0	0	0.00	0	0.00
O2C	US	12	0	0.00	1	8.33	0	0	0.00	0	0.00
SHL	BR	3	0	0.00	3	100.00	0	0	0.00	0	0.00
QMK	BR	298	0	0.00	9	3.02	0	0	0.00	0	0.00

References

1. RIR Statistics. https://www.nro.net/about/rirs/statistics/
2. Alieyan, K., Kadhum, M.M., Anbar, M., Rehman, S.U., Alajmi, N.K.: An overview of DDoS attacks based on DNS. In: 2016 International Conference on Information and Communication Technology Convergence (ICTC), pp. 276–280. IEEE (2016)
3. AMSIX: Pricing — AMS-IX Amsterdam (2023). https://www.ams-ix.net/ams/pricing
4. Antonakakis, M., et al.: Understanding the Mirai botnet. In: 26th {USENIX} security symposium ({USENIX} Security 17), pp. 1093–1110 (2017)
5. Bailey, M., Dittrich, D., Kenneally, E., Maughan, D.: The Menlo report. IEEE Secur. Priv. **10**(2), 71–75 (2012)
6. CAIDA: AS Rank. https://asrank.caida.org/
7. Cymru, T.: network-security-templates/README.md at master · team-cymru/network-security-templates · GitHub (2022)
8. DE-CIX: Blackholing - Fight DDoS attacks effectively. https://de-cix.net/en/services/blackholing
9. Dietzel, C., Wichtlhuber, M., Smaragdakis, G., Feldmann, A.: Stellar: network attack mitigation using advanced blackholing. In: Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies. CoNEXT 2018, pp. 152–164, New York, NY, USA. Association for Computing Machinery (2018). https://doi.org/10.1145/3281411.3281413, https://doi.org/10.1145/3281411.3281413
10. Dnsfilter: Beyond Hackers in Hoodies: DNSFilter Mid-Year Cybersecurity Review (2022)
11. Equinix: Remotely Triggered Black Hole. https://docs.equinix.com/en-us/Content/Interconnection/IX/IX-rtbh-guide.htm
12. Giotsas, V., Smaragdakis, G., Dietzel, C., Richter, P., Feldmann, A., Berger, A.: Inferring BGP blackholing activity in the internet. In: Proceedings of the 2017 Internet Measurement Conference. IMC 2017, New York, NY, USA, pp. 1–14. Association for Computing Machinery (2017). https://doi.org/10.1145/3131365.3131379
13. Jonker, M., Pras, A., Dainotti, A., Sperotto, A.: A first joint look at DoS attacks and BGP blackholing in the wild. In: Proceedings of the Internet Measurement Conference 2018, pp. 457–463 (2018)
14. Jonker, M., Sperotto, A.: Measuring exposure in DDoS protection services. In: 2017 13th International Conference on Network and Service Management (CNSM), pp. 1–9. IEEE (2017)
15. Kopp, D., Dietzel, C., Hohlfeld, O.: DDoS never dies? An IXP perspective on DDoS amplification attacks. In: Hohlfeld, O., Lutu, A., Levin, D. (eds.) PAM 2021. LNCS, vol. 12671, pp. 284–301. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72582-2_17
16. Krämer, L., et al.: AmpPot: monitoring and defending against amplification DDoS attacks. In: Proceedings of the 18th International Symposium Research in Attacks, Intrusions, and Defenses, pp. 615–636 (2015)
17. Kristoff, J.: An Internet-wide BGP RTBH service. Technical report (June 2015). https://www.iab.org/wp-content/IAB-uploads/2015/04/CARIS_2015_submission_20.pdf
18. Krupp, J., Backes, M., Rossow, C.: Identifying the scan and attack infrastructures behind amplification DDoS attacks. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1426–1437 (2016)

19. Lone, Q., Frik, A., Luckie, M., Korczynski, M., van Eeten, M., Gañán, C.: Deployment of source address validation by network operators: a randomized control trial. In: Proceedings of the 43rd IEEE Symposium on Security and Privacy (S&P 2022) (2022)
20. Mirkovic, J., Reiher, P.: A taxonomy of DDoS attack and DDoS defense mechanisms. ACM SIGCOMM Comput. Commun. Rev. **34**(2), 39–53 (2004)
21. Nawrocki, M., Blendin, J., Dietzel, C., Schmidt, T.C., Wählisch, M.: Down the black hole: dismantling operational practices of BGP blackholing at IXPs. In: Proceedings of the Internet Measurement Conference, pp. 435–448 (2019)
22. Srivastava, A., Gupta, B.B., Tyagi, A., Sharma, A., Mishra, A.: A recent survey on DDoS attacks and defense mechanisms. In: Nagamalai, D., Renault, E., Dhanuskodi, M. (eds.) PDCTA 2011. CCIS, vol. 203, pp. 570–580. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24037-9_57
23. Team Cymru: Unwanted traffic removal service. https://www.team-cymru.com/ddos-mitigation-services
24. Team Cymru: UTRS Peering Guide. https://github.com/team-cymru/network-security-templates/blob/master/UTRS-Peering-Guide/README.md
25. Turk, D.: Configuring BGP to Block Denial-of-Service Attacks. RFC 3882 (2004). https://doi.org/10.17487/RFC3882, https://rfc-editor.org/rfc/rfc3882.txt
26. Zhauniarovich, Y., Dodia, P.: Sorting the garbage: filtering out DRDoS amplification traffic in ISP networks. In: Proceedings of the IEEE Conference on Network Softwarization, pp. 142–150 (2019)