

Who's got my back? Measuring the adoption of an internet-wide BGP RTBH Service

Anghel, Radu; Zhauniarovich, Yury; Gañán, Carlos

DOI

[10.1145/3639029](https://doi.org/10.1145/3639029)

Publication date

2024

Document Version

Final published version

Published in

Proceedings of the ACM on Measurement and Analysis of Computing Systems

Citation (APA)

Anghel, R., Zhauniarovich, Y., & Gañán, C. (2024). Who's got my back? Measuring the adoption of an internet-wide BGP RTBH Service. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 8(1), 1-25. Article 3. <https://doi.org/10.1145/3639029>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Who's Got My Back? Measuring the Adoption of an Internet-wide BGP RTBH Service

RADU ANGHEL, Delft University of Technology, The Netherlands

YURY ZHAUNIAROVICH, Delft University of Technology, The Netherlands

CARLOS GAÑÁN, Delft University of Technology, The Netherlands

Distributed Denial-of-Service (DDoS) attacks continue to threaten the availability of Internet-based services. While countermeasures exist to decrease the impact of these attacks, not all operators have the resources or knowledge to deploy them. Alternatively, anti-DDoS services such as DDoS clearing houses and blackholing have emerged. Unwanted Traffic Removal Service (UTRS), being one of the oldest community-based anti-DDoS services, has become a global free collaborative service that aims at mitigating major DDoS attacks through the Border Gateway Protocol (BGP). Once the BGP session with UTRS is established, UTRS members can advertise part of the prefixes belonging to their AS to UTRS. UTRS will forward them to all other participants, who, in turn, should start blocking traffic to the advertised IP addresses.

In this paper, we develop and evaluate a methodology to automatically detect UTRS participation in the wild. To this end, we deploy a measurement infrastructure and devise a methodology to detect UTRS-based traffic blocking. Using this methodology, we conducted a longitudinal analysis of UTRS participants over ten weeks. Our results show that at any point in time, there were 562 participants, including multihomed, stub, transit, and IXP ASes. Moreover, we surveyed 245 network operators to understand why they would (not) join UTRS. Results show that threat and coping appraisal significantly influence the intention to participate in UTRS.

CCS Concepts: • **Security and privacy** → **Denial-of-service attacks; Security services**; • **Networks** → **Network measurement**.

Additional Key Words and Phrases: BGP, RTBH, DDoS, UTRS, Internet measurements

ACM Reference Format:

Radu Anghel, Yury Zhauniarovich, and Carlos Gañán. 2024. Who's Got My Back? Measuring the Adoption of an Internet-wide BGP RTBH Service. *Proc. ACM Meas. Anal. Comput. Syst.* 8, 1, Article 3 (March 2024), 25 pages. <https://doi.org/10.1145/3639029>

1 INTRODUCTION

As the intensity, capacity, and duration of Denial of Service (DoS) attacks increase, network operators constantly look for countermeasures [26]. There are many DDoS mitigation methods, each having different levels of complexity, cost, and efficiency [49, 50]. The most common are blackholing, clean pipe, content delivery networks (CDN) attack dilution, and antiDDoS proxy. While antiDDoS proxy and CDN attack dilution methods are widespread for specific applications (e.g., TCP- or UDP-based), clean pipe and blackholing are the only ones that work for all application types. Given that clean pipe increases the loading times (latency) and is complex to deploy, blackholing is more popular among network operators [13].

Authors' addresses: Radu Anghel, R.Anghel@tudelft.nl, Delft University of Technology, The Netherlands; Yury Zhauniarovich, Y.Zhauniarovich@tudelft.nl, Delft University of Technology, The Netherlands; Carlos Gañán, C.HernandezGanan@tudelft.nl, Delft University of Technology, The Netherlands.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

ACM 2476-1249/2024/3-ART3

<https://doi.org/10.1145/3639029>

With blackholing, all traffic, covering both legitimate and attack packets, is sent into a black hole, or null route. In this context, Remotely Triggered Black Hole (RTBH) emerges as a commonly used technique that provides the ability to dynamically announce attack destinations for which undesirable traffic should be dropped before it enters a protected network [50]. At its core, RTBH leverages the Border Gateway Protocol (BGP) to reroute attack traffic to places that minimize harm, which usually entails dropping the traffic.

When only one network operator removes incoming DDoS traffic, the attack is only partially mitigated and the garbage traffic continues to impact network operators as it traverses the Internet. This reduces their bandwidth and affects their customers. At the same time, if multiple network operators collaborate, the attack may be stopped closer to its source, and more of the path it would have taken is protected. To facilitate this collaboration at no cost, *Unwanted Traffic Removal Service* (UTRS) was launched as a free RTBH solution in 2014. This service allows its members to announce some IP addresses they own, the traffic to which should be dropped. To the best of our knowledge, UTRS is currently the only *global free-to-participate* RTBH initiative, so network operators who do not have access to ISPs offering RTBH or do not want to incur additional costs can still participate in it, and get yet another anti-DDoS tool in their arsenal.

Unfortunately, this collaborative RTBH suffers from network effects [4], i.e., the effectiveness of the blackholing increases as more network operators work together. Therefore, it becomes essential to know how many network operators participate in UTRS, who they are, and what their characteristics are in order to be able to evaluate how effective it is. This information allows for a comprehensive assessment of the collaborative effort's impact on DDoS attack mitigation. By understanding how many network operators are actively engaged in RBTH and who they are, researchers and practitioners can gauge the scope and reach of the system. Moreover, gaining insights into the characteristics of participating network operators is crucial for evaluating the overall effectiveness of RBTH. Factors such as their geographical distribution, network infrastructure capabilities, and experience in handling DDoS attacks can significantly influence the collective response to threats. This knowledge empowers network operators to make informed decisions about joining the RBTH service.

Therefore, the main research question that we set out is: *How can we identify network operators that are actively UTRS participants?* In this work, we answer this question and characterize UTRS participants by looking at factors like network type, and the institutional environment of the country where the network operator is located. In addition to that, we ran a large-scale survey study among network operators to identify the determinants that lead them to participate in UTRS. Thus, the contributions of this paper are the following:

- We present the first experimental design to identify UTRS participants based on large-scale measurements.
- We perform a longitudinal analysis over a 10-week period identifying at least 562 UTRS members (conservative estimation) that actively participate in the service. We characterize these operators based on their network size, type and operational region.
- We compare the results of our measurements with the results collected using RIPE Atlas probes. This helps us to validate the main measurement methodology.
- We conduct a large-scale online survey across members and non-members of UTRS to validate the results of the active measurements. We received 245 valid responses from both UTRS participants and non-participants. Out of them, 58 out of 59 ASes (98%), who were identified by our approach as UTRS participants, have confirmed their participation in UTRS.
- We identify the behavioral determinants driving UTRS participation using Protection Motivation Theory (PMT). Our results show that perceived vulnerability, perceived severity, response efficacy, response cost, as well as social norms, influence the intention to participate in UTRS.

2 BACKGROUND

BGP [37] is a path-vector routing protocol responsible for ensuring the interconnectivity of ASes (also referred to as *ASNs* - AS Numbers) over the Internet by exchanging information about external routing. In simple terms, BGP is a mechanism that is used to transport information about available traffic routes from one section of the Internet to another. The intricacy of connections, rules, and economics spurred the need for similarly sophisticated and fine-grained routing strategies as the Internet expanded.

One such BGP extension defines *RTBH*. RTBH is one of the oldest methods used to mitigate DDoS attacks, first documented in 2004 [44]. The RTBH works by using a previous arrangement between the victim ISP and the upstream ISPs, peers, or Internet eXchanges (IX). The victim's ISP, using BGP, advertises the IP under attack, sometimes also using a previously agreed BGP community to its upstreams, peers, or IX route servers. Upon receiving such an announcement, they start discarding packets to that destination (null route, blackhole). This action has the effect of blocking all traffic toward the victim's IP, which also completes the attack because the victim is deemed offline.

Granular control over what is blocked can be achieved by using the FlowSpec format [31]. This format allows specifying fine-grained rules similar to ACL/firewall ones. These rules allow peers to perform more selective blocking of the traffic. For instance, in case of a reflection DDoS attack using NTP, the AS can ask to block only traffic coming from source port 123 (UDP/NTP) to the attacked host.

UTRS is a community project of Team Cymru [40]. Simply put, UTRS is an RTBH operated by a trusted third party. Any AS can become a UTRS member and use the service to announce IPs, the traffic to which should be blocked. Figure 1 explains this process. During Step 1, a UTRS member announces an IP to UTRS. Members can announce networks under their control up to /25 and /49 size for IPv4 and IPv6 correspondingly. During Step 2, UTRS distributes these announcements among all the members. UTRS also supports BGP FlowSpec [31]. These announcements through UTRS reach all members, i.e., they all receive information about which IPs should be blackholed. However, UTRS does not force the members to do that. Thus, it is possible that some UTRS members do receive the announcements but do not block the corresponding IP addresses. We call the UTRS members that actively block the traffic to the announced IPs *Active UTRS Members*. Note the collaboration within UTRS allows stopping the DDoS traffic closer to its source [28, 40]; thus, preventing it not only from reaching the destination network but also from congesting the networks on the path to the victim.

3 OVERVIEW

3.1 System Setup

Figure 1 presents a high-level overview of our testbed. To perform our measurements, we used an Autonomous System (Test AS) and a /24 IPv4 range under our control. The Test AS participates in UTRS, but for measurement purposes, it does not actively block packets. UTRS validates authority over the AS by getting confirmation from registered contacts in the RIR (Regional Internet Registry) database for that AS. Further verification is performed automatically by the system when routes are sent to UTRS: the IP(s) must be registered to the participating AS in order to be accepted and propagated to the other participants.

We assigned three IPv4 addresses from the /24 in our Test AS to a virtual machine (VM): A(nother), B(locked) and C(ontrol). We use the A(nother) IP (*A-IP*) to identify pingable IPv4 addresses in each AS visible in the Global Routing Table. We use a separate IP address for scanning to minimize the effects of potential automatic blockings. The B(locked) IP (*B-IP*) is the one that we have announced to UTRS for blackholing. In theory, the other participating AS should block




Fig. 1. UTRS Announcement and System setup

traffic to this IP address after receiving the announcement. The last C(ontrol) IP address (C-IP) is used for control purposes to confirm the connectivity between the Test AS and the tested AS.

3.2 Measurement Methodology

Figure 2 outlines the methodology we used in this paper to identify ASNs actively participating in UTRS. For the sake of clarity, in this work, we call ASNs discovered using our approach as **Detected Active UTRS Participants** (or **UTRS Participants** for short) in order to distinguish them from the **UTRS Members** (actual members of UTRS). Note that there are no checks or penalties for a UTRS member for not blocking the traffic to the announced IP. It is possible that some participating ASNs are only receiving the UTRS-related route updates but do not block traffic to those IPs (e.g., our Test AS is a UTRS member but is not an active UTRS participant – it does not blackhole announced IPs). Unfortunately, our approach is not able to detect such cases. Instead, we identify only ASes actively participating in UTRS, i.e., the ones that perform blackholing. Hereafter, we describe each step of our approach in detail.




Fig. 2. UTRS participants identification methodology

Step 1: Announcing B-IP to UTRS. Before launching an experiment, we announce *B-IP* to UTRS. The *B-IP* is announced only once before the start of the measurements and withdrawn when they are done, so Route Flap Dampening [48] is not affecting the measurements. As a result, all networks of ASes participating in UTRS should start blocking traffic to this IP.

Step 2: Identifying a Pingable IP in each AS. During the second step, we look for a pingable IP address in every AS that we use in our further measurements. From the latest Routing Information Base (RIB) available at the University of Oregon’s RouteViews project [45], we gather a list of maximum 10 prefixes for each AS visible in the Global Routing Table. Then, using the ZMap’s [16] `icmp_echoscan` module, we randomly ping IPs in these networks, stopping when a pingable IP is found. For this scan, we use *A-IP* as a source to reduce negative effects on our further measurements (e.g., some network firewalls may start blocking our *A-IP* after some threshold is reached, but this should not affect measurements from *B-IP* and *C-IP*). As a result, we get a list of the ASNs and the corresponding pingable IP addresses. Note that we use our own ZMap scan results instead of Censys [15] or USC ISI [17] hitlists because, for our measurements, it is crucial to have the freshest possible data, that would still require us to ping the IPs from these lists.

Step 3: Pinging IPs identified in Step 1 from B-IP and C-IP. We ping each IP address identified in the previous step from *B-IP* and *C-IP*. We chose to use ICMP due to its high simplicity and

reachability, with a 98.8% success rate compared to 88.7% for TCP and only 14.3% for UDP, as found by Bano et al. [8]. Note that at *Step 1*, *B-IP* has been announced to UTRS. Because of this, all traffic to *B-IP*, including ICMP echo replies, must be blocked by active UTRS members. Thus, our *B-IP* node will not receive ICMP echo replies from IP addresses belonging to networks actively participating in UTRS. Additionally, *B-IP* will not receive replies if the traffic to it passes through a network of an active UTRS member.

C-IP is not announced to UTRS. Thus, all the traffic to this IP should not be blocked by UTRS members so that we will get ICMP echo replies. For instance, in [Figure 1](#), we should get ICMP echo replies from pingable IP addresses in the AS2, AS3, AS4, AS6, and AS7 networks if we ping them from the *C-IP* source but will not receive anything when pinging them from *B-IP*.

We ping IPs from *B-IP* and *C-IP* almost simultaneously: right after pinging from the *B-IP* node, we launch another ping from *C-IP*. It is highly unlikely that our whole Test AS network would be blocked by a firewall or an Intrusion Detection System (IDS) in such a short time. Moreover, to each pingable IP address, we send three packets from *B-IP* and *C-IP*. Such a low number should not result in blocking our network as well.

Step 4: Comparing Ping Results from B-IP and C-IP. At this step, we compare the replies collected from each ping at the previous step. There are four possible outcomes:

Case BCxx ($\#Pkt_{B-IP} > 0, \#Pkt_{C-IP} > 0$): Pings from both *B-IP* and *C-IP* to the pingable IP address get replies. We consider the corresponding AS as one that does not participate in UTRS.

Case BC0x ($\#Pkt_{B-IP} = 0, \#Pkt_{C-IP} > 0$): Pings from *C-IP* receive a reply, while those sourced from *B-IP* do not. We consider the corresponding AS as affiliated with UTRS (the AS can be either an active UTRS participant or routes to it pass through a network of an active UTRS participant).

Case BCx0 ($\#Pkt_{B-IP} > 0, \#Pkt_{C-IP} = 0$): Pings from *B-IP* receive a reply, while from *C-IP* do not. Ideally, this should not happen, but we observe such cases. Their presence can be explained by connectivity issues, ICMP rate-limiting, and multipath routing (packets take different routes from/to the destination: one blocks, while another does not).

Case BC00 ($\#Pkt_{B-IP} = 0, \#Pkt_{C-IP} = 0$): Pings from both *B-IP* and *C-IP* do not receive replies. This result is unexpected as we have got a reply to *A-IP* before. However, this may be caused by several reasons. For instance, this could be due to the connectivity problem with/for the corresponding AS, ICMP rate-limiting, or a firewall blocking packets.

Step 5: Traceroute IPs not reachable from B-IP. We consider the ASes from *Case BC0x*, identified during the previous step, as active UTRS participant candidates: the traffic from this candidate AS to the IP announced to UTRS (*B-IP*) is blocked. This can happen for two reasons: either the AS actively participates in UTRS itself, or the traffic to it passes through a network of an active UTRS member. For instance, in [Figure 1](#), AS2, AS4 and AS6 are active UTRS members, while AS7 is not. However, the traffic to AS7 passes through the network of AS4, which performs the blackholing. Therefore, from our Test AS point of view, AS7 is also considered as the AS participating in UTRS. AS3 is a particular case. The traffic to it passes through ASes participating in UTRS, but itself, it is also an active UTRS participant.

To determine which candidate ASes are active UTRS participants, we rely on the traceroutes to the pingable IP addresses corresponding to the ASes potentially participating in UTRS. Our initial intuition was the following. If a candidate AS is an active UTRS participant, then the traceroute sourced from *B-IP* to a pingable IP would stop at a router that belongs to this AS. Therefore, comparing the traceroutes from *B-IP* and *C-IP* to the pingable IP, we should be able to distinguish active UTRS participants from the upstream blocking ASes. Unfortunately, the initial analysis of the collected data using the method described in the previous step showed the impracticability of

our naïve approach. There are several issues with it. First, traceroutes from *B-IP* and *C-IP*, even being launched at the same exact moment may go forward through different paths, for instance, in the case of Equal-Cost Multipath (rfc2992) [23]. Second, the backward path of a reply from a router may not be the same as the forward path of the corresponding request. A reply may go through a network participating in UTRS, thus, forcing us to falsely mark the router's AS as participating in UTRS. These peculiarities make our naïve approach unusable.

Step 6: Prune the Data using Traceroutes. To overcome these limitations of the naïve approach described in the previous step, we have developed a new probabilistic graph-based approach for detecting the ASes participating in UTRS. Our approach relies only on the traceroutes collected from non-blocked *C-IP* and ping data from *B-IP* and *C-IP*. It takes into consideration all traceroute measurements from *C-IP* simultaneously and combines them with the information obtained at Step 3. This provides a broader view of the relationships between the ASes, allowing us to better filter out the ASes not participating in UTRS, however, making our approach probabilistic. In the following section, we consider our algorithm in detail.

3.3 Data Pruning Algorithm

The data pruning process consists of two phases. During the first phase, we build a directed graph representing *Path of transited ASes*. During the second phase, we analyze each node in the graph, selecting those who, with *high probability*, are active UTRS participants.

Phase 1: Bulding the Graph of Transited ASes. This phase includes several stages. First, we take the candidate list of active UTRS participants obtained at Step 4 and the traceroutes from *C-IP* to the pingable IPs in these ASes. During the second stage, we analyze each traceroute to a pingable IP. In particular, we take the IP address of each hop in a traceroute and map it to the corresponding AS number (ASN). For this data enrichment, we use the *pyasn* [21] Python library with the latest route information database from the RouteViews project [45] available at the time of measurement. If *pyasn* cannot detect the ASN (e.g., if the IP is in a private or not announced network - Internet Exchange ranges are usually not announced to the Internet to protect them from attacks), we ignore these nodes. Thus, we obtain a sequence of ASNs representing a path from our Test AS to the ASN corresponding to the pingable IP. During the third stage, we remove multiple occurrences of the same ASN (in this case, the sequence represents several consecutive routers belonging to the same AS), i.e., self-loops in the graph.

Then, we add the nodes from a sequence to our graph. If a node is in our candidate list of active UTRS participants obtained at Step 4, we mark it as *blocking*, otherwise, we mark it as *non-blocking*. In Figure 1, blocking nodes are red (AS2, AS3, AS4, AS6 and AS7), while non-blocking (AS1 and AS5) are blue. The green node, Test AS, is our autonomous system. It is considered non-blocking in further analysis. The AS graph in Figure 1 resembles the data-plane view from the Test AS to the tested ASes representing the paths the packets actually take to travel to their destinations.

Phase 2: Nodes Analysis. During the second phase, we analyze each node in the graph and its predecessors. Algorithm 1 provides the pseudocode of the graph analysis algorithm.

We iterate over each node in the graph (Line 3) and analyze whether it is blocking or non-blocking. If the node is blocking (Line 5), then we check if any of the node's predecessors is non-blocking (see Line 6). If this is the case, we add this node to the set of *active UTRS participants*. In Figure 1, AS2, AS4 and AS6 belong to this set. Note that we can make this conclusion with a very high probability. If all predecessors are blocking (e.g., AS3 or AS7), then we cannot make such a conclusion: *the corresponding AS may be an active UTRS participant (AS3) but may also block the traffic due to upstream blocking nodes (AS7)*. Thus, AS3 is a False Negative.

Algorithm 1: Graph analysis algorithm

```

In :  $G$  - AS path graph built at Stage 1
Out:  $utrs\_ptcpnts$  - a set of ASNs corresponding to active UTRS participants,
       $odd\_asns$  - a set of ASNs corresponding to Odd ASes

1  $utrs\_ptcpnts \leftarrow \emptyset$ ;
2  $odd\_asns \leftarrow \emptyset$ ;
3 for  $node$  in  $G$  do
4    $pds \leftarrow G.predecessors(node)$ ;
5   if  $blocking(node)$  then
6     if  $any(!blocking(pds))$  then
7        $utrs\_ptcpnts.add(node)$ ;
8     end
9   else
10    if  $all(blocking(pds))$  then
11       $odd\_asns.add(node)$ ;
12    end
13  end
14 end

```

If a node is non-blocking (see [Line 9](#)), then we check if all the node's predecessors are blocking ([Line 10](#)). If yes, we add the node to the *Odd ASes* set (the traffic reaches this node although all upstream nodes block the traffic). For instance, in [Figure 1](#), AS5 represents a member of this set. As you can see, the traffic to this node comes from the ASes that actively participate in UTRS; therefore, ideally, the traffic from our Test AS should not reach AS5. Additional routes to node AS5 or routing misconfiguration might explain the appearance of these cases (they represent only 1% of the cases).

It is possible to prune the list of active UTRS participants further by considering the results of several independent measurements. However, this approach has several issues. First, our measurements are distinct in time (one week). During this period, the list of UTRS members may change (new members can join UTRS, or some may stop participating). If we run the experiments more often, the results may also be incorrect (e.g., some networks may start blocking the packets from our Test AS). Second, it is not clear what threshold to choose. In order to get the threshold value, we need ground truth data which, unfortunately, is not available. We leave investigating these issues as future work. Within the scope of this paper, we provide the results considering each measurement individually and report a conservative estimate of UTRS participants number by selecting the ones who appear in all our measurements (see [Section 4.3](#)).

3.4 RIPE Atlas Experiment

Due to the probabilistic nature of our method, a single measurement using our approach does not produce a complete and entirely accurate list of UTRS participants. Authors in [\[29, 36\]](#) observed that some routers have stringent ICMP rate limiting rules. As we issue a non-negligible number of ICMP requests during the Main Experiment, it may occur that some ISPs will block the pings from our AS after some time, which may negatively affect the results of our analysis. Due to Team Cymru's [\[40\]](#) refusal to share the list of UTRS members to verify our findings directly, we add an additional experimental setup to validate our results. To achieve this goal, we use the RIPE Atlas platform [\[1\]](#). This platform allows executing network measurements, including pings and traceroutes, from probes spread out all over the world.

We employ this platform to run an experiment similar to the one described in [Section 3.2](#). However, in this experiment, we do the measurements in the reverse direction, i.e., we ping and traceroute our *B-IP* and *C-IP* from the Atlas probes. Thus, contrary to our Main Experiment, we do not expect the corresponding probe IP addresses to be blocked as the number of ICMP echo

packets from a probe is negligible. Moreover, this experiment could help us to discover more UTRS participants. Because the measurements are done from the other side, we may uncover UTRS participants that are considered as upstream blocked in our main experiment (e.g., AS3 in Figure 1). We adapted our methodology and algorithm described in Section 3.2 and Section 3.3 to this setup:

Step 1: Obtain a list of probes from the Atlas API filtered by status=Connected and tagged as system-ipv4-works.

Step 2: Run ping measurements towards *B-IP* and *C-IP*. Due to the limitations of the Atlas platform, we split the probes in batches to be able to run the measurement with all available probes.

Step 3: Compare ping results:

Case BCxx: Pings to both *B-IP* and *C-IP* succeed - non blocking;

Case BC0x: Ping to *B-IP* fails while to *C-IP* succeeds - blocking;

Case BCx0: Ping to *C-IP* fails while to *B-IP* succeeds - should not happen in theory;

Case BC00: Both pings to *B-IP* and *C-IP* fail - there is a connectivity problem, firewall, packet loss.

Step 4: Trace routes to *B-IP* (not used) and *C-IP* (used) for the *Case BC0x* probes.

Step 5: Process traceroutes using an adapted Algorithm 1.

4 MEASUREMENT RESULTS

We ran our measurements once per week for *ten* consecutive weeks in March-May 2022. Each week we performed two measurements, namely ‘Main Experiment (ME)’ and ‘RIPE Atlas Experiment (RAE)’. In this section, we analyze the results of each experiment separately, and compare them as validation of the main measurement methodology.

4.1 Individual Measurements Results

Table 1 shows the results of the individual measurements for our experiments. It is split into two parts headed ME and RAE corresponding to each measurement. These two parts have a similar structure: firstly, we present the initial population numbers for both experiments, then we report the number of ASes of depending on ping reachability, and finally, we provide the number of UTRS Participants identified using our algorithm for each individual week.

Table 1. Number of detected ASes per measurement: ME - Main Experiment; RAE - RIPE Atlas Experiment; UP - UTRS Participants; OA - Odd ASes

| Number of | Measurement | | | | | | | | | |
|------------------------------------|-------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| Main Experiment (ME) | | | | | | | | | | |
| Active ASes | 72768 | 72740 | 72771 | 72788 | 72854 | 72908 | 72973 | 72960 | 73060 | 73038 |
| ASes w/ Pingable IPs | 68354 | 68288 | 68272 | 68309 | 68402 | 68500 | 68583 | 68503 | 68669 | 68600 |
| <i>BCxx</i> _{ME} | 65712 | 65409 | 65614 | 66031 | 66295 | 66179 | 65978 | 66008 | 66157 | 66414 |
| <i>BC0x</i> _{ME} | 2668 | 2778 | 2783 | 2188 | 2207 | 2237 | 2485 | 2476 | 2423 | 2592 |
| <i>BCx0</i> _{ME} | 44 | 53 | 61 | 50 | 74 | 57 | 65 | 57 | 41 | 48 |
| <i>BC00</i> _{ME} | 85 | 63 | 68 | 91 | 73 | 94 | 82 | 54 | 60 | 105 |
| $\bar{U}P_{ME}$ | 1498 | 1529 | 1570 | 1217 | 1252 | 1321 | 1397 | 1356 | 1426 | 1470 |
| <i>OA</i> _{ME} | 14 | 14 | 13 | 11 | 11 | 13 | 14 | 11 | 14 | 14 |
| RIPE Atlas Experiment (RAE) | | | | | | | | | | |
| Probes | 11088 | 11053 | 11051 | 11083 | 11084 | 11098 | 11117 | 11135 | 11168 | 11157 |
| ASes w/ Probes | 3483 | 3483 | 3493 | 3502 | 3500 | 3498 | 3489 | 3505 | 3519 | 3505 |
| <i>BCxx</i> _{RAE} | 3379 | 3384 | 3388 | 3400 | 3397 | 3395 | 3385 | 3399 | 3403 | 3391 |
| <i>BC0x</i> _{RAE} | 110 | 108 | 113 | 108 | 103 | 105 | 109 | 114 | 123 | 126 |
| <i>BCx0</i> _{RAE} | 5 | 3 | 3 | 4 | 5 | 1 | 2 | 1 | 2 | 4 |
| <i>BC00</i> _{RAE} | 21 | 19 | 24 | 19 | 22 | 19 | 20 | 20 | 20 | 21 |
| $\bar{U}P_{RAE}$ | 90 | 89 | 94 | 91 | 88 | 91 | 94 | 97 | 100 | 101 |
| <i>OA</i> _{RAE} | 2 | 3 | 2 | 2 | 3 | 3 | 2 | 2 | 1 | 3 |

As we explained in [Section 3.2](#), we start our Main Experiment (ME) by collecting a list of *Active ASes* and the corresponding 10 network prefixes from the RIBs provided by the RouteViews project [45]. The “Active ASes” row in [Table 1](#) provides the number of *Active ASes* we got from the RIB files for a particular measurement. This number is pretty stable during these weeks and varies in the range 72,740 – 73,060 ASes, increasing as more ASes are assigned by RIRs and put in use. Out of these ASes, each week, on average, we discovered 68,448 Autonomous Systems with Pingable IPv4 addresses (see the “ASes w/ Pingable IPs” row for exact values per each week).

For the RIPE Atlas Experiment (RAE), we report correspondingly the number of probes (“Probes”) and ASes (“ASes w/ Probes”) where these probes are located. In the RIPE Atlas network, there are about 11,000 such probes in approximately 3,500 distinct ASes. The coverage of ASes in RAE is considerably lower than in ME. Comparing the numbers in [Table 1](#), we can see that in ME, pingable IPs are discovered in 68,448 ASes, while in RAE, probes are only found in 3,498 ASes on average. Thus, *Main Experiment covers almost 20 times more ASes than RIPE Atlas Experiment*.

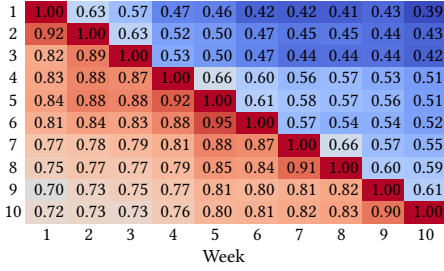
The rows “ BCx_{ME} ”, “ $BC0_{ME}$ ”, “ $BCx0_{ME}$ ”, and “ $BC00_{ME}$ ” in [Table 1](#) show the number of ASes per each case described in [Step 4](#) in [Section 3.2](#). The majority of the ASes with pingable IPs do not participate in UTRS (on average 96.39% across all weeks). There could be multiple reasons for that: network operators do not know about the UTRS project; they do not trust Team Cymru [40]; ISPs do not see an immediate benefit in its adoption, or they use a different RTBH service, e.g., their own or provided by an Internet EXchange Point (IXP). In [Section 6](#), we explore these factors.

Theoretically, the ASes from the $BCx0$ category should not be present in our results because this means that the announced B -IP gets the response, while C -IP does not. This situation is possibly explained by the target network starts blocking us after receiving the pings from our A - and B -IPs, and our C -IP does not get the response. Moreover, such situations are also possible due to Internet volatility due to connectivity issues or multipath routing. However, the number of such ASes is low, constituting only 0.08% of all ASes with pingable IPs. We can consider the corresponding ASes as the ones that do not participate in UTRS because B -IP gets replies. Similarly, the presence of the $BC00$ instances can be explained by the blocking happening after the initial ping from A -IP. The portion of $BC00$ ASes is larger than the $BCx0$ ones and is equal, on average, to 0.1%.

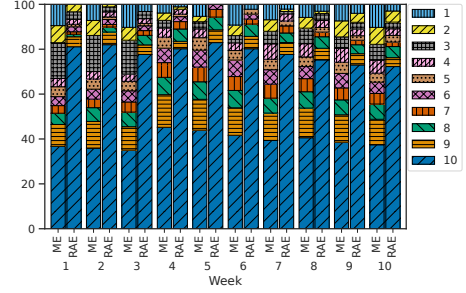
Similarly to ME results, the majority of the ASes, where RIPE Atlas probes are located, do not participate in UTRS (on average 96.98% across all weeks). Out of 3,498 checked ASes, we were able to find, on average, only 112 of them that potentially participate in UTRS (see the “ $BC0_{RAE}$ ” row in the lower part of [Table 1](#)). We also found $BCx0$ (0.09%) and $BC00$ (0.59%) cases in our RIPE Atlas data. This shows that these cases are a universal phenomenon and are not attributed only to our Test AS. Rather, they exist due to the flaky nature of the Internet. Note that the portion of $BC00$ cases in RAE is higher than in ME. This may be due to traffic blocking from the RIPE Atlas probes or because we do not test the connection to our Test AS from RIPE Atlas probes (similar to pings from our A -IP in our ME).

The values in the “ $BC0_{ME}$ ” and “ $BC0_{RAE}$ ” rows show the number of ASes that block the traffic to B -IP while allowing the traffic to C -IP. These ASes are either active UTRS participants, or the traffic to our Test AS from them passes through networks of ASes actively participating in UTRS. According to our measurements, these values represent the upper bound of the number of active UTRS participants. According to our data, 4.08% of ASes with pingable IPs and 3.2% of distinct ASes with probes are active UTRS participants at most. We feed these sets of ASes together with the traceroute data to our data pruning algorithm. As its output, we obtain two sets: the first stores UTRS Participants (UP), and the second contains Odd ASes (OA).

The “ UP_{ME} ” row in [Table 1](#) shows the number of active UTRS participants for each week in ME. Using this approach and the data, we have identified between 1,217 up to 1,570 ASes that actively participate in UTRS, with an average of 1,404 (1.93% of all active ASes). Unfortunately,



(a) Jaccard similarity of UTRS participants (ME results – above the diagonal, RAE – below)



(b) Occurrence frequency of UTRS participants in ME and RAE

Fig. 3. Robustness of the results

Team Cymru does not reveal the list of UTRS members as well as their exact number (although it did this in the past). Therefore, we cannot verify our findings with the ground truth data. At the same time, their webpage [40] says that their community contains 1300+ network operators.

At the same time, using the data from RAE, we detect from 88 to 101 active UTRS participants (“ UP_{RAE} ”). On average, 2.67% of distinct ASes with RIPE Atlas Probes are active UTRS participants. This percentage is higher than for ME where 1.93% of all active ASes belong to active UTRS participants. There could be several explanations for this phenomenon. We hypothesize that the networks that host RIPE Atlas probes are managed by more Internet-savvy administrators. Therefore, there is a higher chance that they would also push to become UTRS members.

Note that the number of detected active UTRS participants is not stable. For instance, between the third and the fourth week, we observe a sharp drop (by 353) in the number of active UTRS participants (from 1,570 to 1,217). Such a rapid drop cannot be explained by natural fluctuations of UTRS member numbers (ASes joining and leaving UTRS) and probably, reflect network issues. The number of UTRS participants detected using the RAE data confirms this (there is no such rapid decrease). We can further prune the results using the data from several measurements (see Section 3.3). However, in order to do this, we would need access to ground truth to estimate the threshold value. In this paper, we consider as UTRS participants only those ASes that consistently appear in our results during all ten weeks.

4.2 Intra-Experiment Results Analysis

In this section, we analyze the sets of active UTRS participants within one experiment, which we identified using our methodology. Our goal is to estimate if we can rely on one measurement to obtain a robust set of UTRS participants. I.e., if we get roughly the same set of UTRS participants for each measurement, then there is no need to do several measurements in order to improve the reliability of the results. To achieve this goal, we calculate the Jaccard similarity between UTRS participant sets for a pair of different measurements. So as Jaccard similarity is a symmetrical metric, Figure 3a presents the results for ME and RAE simultaneously: the ME and RAE data are above and below the diagonal correspondingly. The figure visualizes the results using a heat map: red and blue colors match high and low values of Jaccard similarity correspondingly.

We can draw several conclusions from Figure 3a. First, the closer the measurements in time, the higher the similarity between the UTRS participant sets both for Main and RIPE Atlas Experiments (the values consistently increase the closer to the diagonal). Indeed, for ME, the highest Jaccard similarity (0.66) is between the sets of the fourth and fifth weeks, while the lowest (0.39) is between

the first and the tenth. Similarly, the RAE results show the highest similarity between the fifth and the sixth week (0.95), while the lowest (0.70) is recorded between the sets of the first and the ninth weeks. That shows the community is not yet stable: the members have been joining and leaving (or stopping to block the advertised IPs). This suggests that we should do the measurements more often to get a better picture of dynamics. We consider this as future work.

Second, we can clearly see that the sets obtained using the RAE data are more stable over time than the UTRS participant sets obtained using the data from ME: upper part of [Figure 3a](#) is more bluish, while the lower part is more reddish. This means that the sets of UTRS participants are more similar for RAE than ME. [Figure 3b](#) reflects this observation even more clearly. It shows the proportion of detected UTRS participants occurring a particular number of times in each week's data. For each week, there are two columns: one column for ME and another for RAE.

The proportion of ASes appearing in every week's UTRS participants set is larger in RAE than in ME. Indeed, on average, around 78% (or 73) of UTRS participants are the same across all weeks for RAE, while only 39% (or 550) of ASes are the same for ME. We assume that a much lower percentage in ME may be explained by different UTRS policies for different network prefixes within the same AS. We elaborate on this issue in [Section 7](#).

At the same time, both the RAE and ME data have UTRS participants that appear only once. On average, there are about 7.57% and 1.87% of such ASes in ME and RAE correspondingly. These results confirm that Internet measurements are flaky. Moreover, they also suggest that our probabilistic algorithm still has space for improvement.

During all ten weeks, we observed 3259 unique ASes identified as UTRS participants in ME and 128 in RAE. Out of them, 550 (16.88%) and 73 (57.03%) are met in each ME and RAE measurement correspondingly. We speculate that, with a very high probability, these ASNs are UTRS members because they consistently show in the results. However, to create a final list of UTRS, we first need to find out what UTRS participants identified in both experiments. In [Section 4.3](#), we cross-analyze the results of these experiments and create a final set of UTRS participants.

4.3 Inter-Experiment Results Analysis

In this section, we cross-compare the sets of Autonomous Systems that we obtained in ME and RAE. [Table 2](#) reports the results of this analysis. Compared to [Table 1](#), there are two additional columns: "Unique" reports the size of the union of all corresponding measurement sets, and "Common" shows the size of the intersection.

The first row, " CA_{ME} ", shows the number of ASes checked during ME. The values there correspond to the number of ASes with Pingable IPs in [Table 1](#). During all ten weeks, 70,975 ASes are checked at least once in our ME, out of which 65,364 are the same across all measurements. The " CA_{RAE} " row reports the number of checked ASes during RAE, i.e., the number of ASes with RIPE Atlas probes. Interestingly, even though adding to or removing a node from RIPE Atlas network is not a frequent event, the number of RAE checked ASes fluctuates considerably. During all measurements, we checked 3,815 unique ASes, out of which only 3,073 are met in all measurements.

While we assumed that the coverage of AS during ME was extensive, the RAE data can still augment it. This becomes obvious if we consider the values in the third and fourth rows in [Table 2](#). Indeed, for each measurement, the size of the union of ME and RAE ASes (the values in the " $CA_{ME} \cup CA_{RAE}$ ") is larger than the number of ME Checked ASes (" CA_{ME} "), while the size of the intersection (" $CA_{ME} \cap CA_{RAE}$ ") is smaller than the number of RAE Checked ASes (" CA_{RAE} "). In total, using the data from both experiments, we were able to cover 70,992 distinct ASes, out of which 65,401 are common in all measurements.

This is also visible in the results provided in the second part of [Table 2](#), where we compare the sets of UTRS participants detected during ME and RAE. Indeed, the number of UTRS participants

Table 2. Inter-experiment results: ME - Main Experiment; RAE - RIPE Atlas Experiment; UP - UTRS Participants; CA - Checked ASes

| | Measurement | | | | | | | | | | Unique | Common |
|---|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|------------|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | | |
| <i>Number of Checked ASes (CA)</i> | | | | | | | | | | | | |
| CA_{ME} | 68354 | 68288 | 68272 | 68309 | 68402 | 68500 | 68583 | 68503 | 68669 | 68600 | 70975 | 65364 |
| CA_{RAE} | 3483 | 3483 | 3493 | 3502 | 3500 | 3498 | 3489 | 3505 | 3519 | 3505 | 3815 | 3073 |
| $CA_{ME} \cup CA_{RAE}$ | 68383 | 68314 | 68297 | 68331 | 68423 | 68522 | 68610 | 68526 | 68694 | 68623 | 70992 | 65401 |
| $CA_{ME} \cap CA_{RAE}$ | 3454 | 3457 | 3468 | 3480 | 3479 | 3476 | 3462 | 3482 | 3494 | 3482 | 3798 | 3036 |
| <i>Number of UTRS Participants (UP)</i> | | | | | | | | | | | | |
| UP_{ME} | 1498 | 1529 | 1570 | 1217 | 1252 | 1321 | 1397 | 1356 | 1426 | 1470 | 3259 | 550 |
| UP_{RAE} | 90 | 89 | 94 | 91 | 88 | 91 | 94 | 97 | 100 | 101 | 128 | 73 |
| $UP_{ME} \cup UP_{RAE}$ | 1521 | 1548 | 1596 | 1238 | 1273 | 1341 | 1412 | 1378 | 1448 | 1494 | 3287 | 579 |
| $UP_{ME} \cap UP_{RAE}$ | 67 | 70 | 68 | 70 | 67 | 71 | 79 | 75 | 78 | 77 | 100 | 44 |
| <i>RAE UTRS Participants not in ME UTRS Participants ($\Delta = UP_{RAE} - UP_{ME}$)</i> | | | | | | | | | | | | |
| Δ | 23 | 19 | 26 | 21 | 21 | 20 | 15 | 22 | 24 | 24 | 62 | 7 |
| $\Delta - CA_{ME}$ | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 |
| $\Delta \cap BCxx_{ME}$ | 15 | 10 | 14 | 13 | 11 | 11 | 9 | 13 | 15 | 13 | 46 | 3 |
| $\Delta \cap BC0x_{ME}$ | 8 | 9 | 10 | 7 | 9 | 8 | 4 | 7 | 5 | 10 | 18 | 3 |
| $\Delta \cap BCx0_{ME}$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| $\Delta \cap BC00_{ME}$ | 0 | 0 | 1 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 2 | 0 |
| Final UTRS Participants | 1506 | 1538 | 1582 | 1225 | 1262 | 1330 | 1403 | 1365 | 1433 | 1481 | 3265 | 562 |

in the ME&RAE sets union ($UP_{ME} \cup UP_{RAE}$) is higher than only in the ME UTRS participants set itself (UP_{ME}). There are 550 and 73 common UTRS participants during all ten weeks detected using the ME and RAE data correspondingly.

We also analyzed the difference between RAE and ME UTRS participant sets for each week. The “ Δ ” row reports the number of ASes detected as UTRS participants using the RAE data that are not detected as those in ME. As you can see, there are about 21 such ASes. That shows that despite higher coverage of ASes during the Main Experiment, RIPE Atlas Experiment is still useful and contributes to the results. There are several reasons why this happens.

Although we have shown that ME covers a considerably larger number of ASes than RAE, it is possible that we are not able to find pingable IPs in some ASes hosting RIPE Atlas probes. Therefore, we are not able to confirm if they are UTRS participants using the ME data. We compared the sets of RAE UTRS participants (UP_{RAE}) and ME Checked ASes (CA_{ME}). The row “ $\Delta - CA_{ME}$ ” reports the difference. As you can see, there is only 1 such unique AS. This AS has only one /24 IPv4 prefix assigned. As this AS is not covered by our ME, we consider it as a UTRS participant only based on the RAE data and add it to the final set of UTRS participants.

Consequently, all other RAE UTRS Participants are also checked during ME, but they are not detected as ME UTRS Participants. To understand why this happened, we get the intersection of these ASes with the sets of ASes corresponding to different ME Ping cases. The four rows with emphasized titles in the third part of Table 2 report these values. Three out of four rows have non-zero values. Let us consider each case in detail.

Intersection with ME BCxx (“ $\Delta \cap BCxx_{ME}$ ”). The ASes from this category were detected as RAE UTRS Participants, while during ME, we received echo replies to both *B-IP* and *C-IP*. Most likely, the corresponding ASes represent *False Positives for the RIPE Atlas experiment* – they are falsely identified as UTRS participants using RAE data.

Intersection with ME BC0x (“ $\Delta \cap BC0x_{ME}$ ”). The ASes from this category were detected as candidates during ME, but were filtered out with our pruning algorithm because there is no link to them from a non-blocking AS. At the same time, during RAE, when we do the measurements in the reverse direction, these ASes were identified as UTRS participants. Thus, the numbers in this

row can be interpreted as *False Negatives for our Main Experiment*. However, these are not all False Negatives. To get the correct number, we should have had probes in all ASes and run measurements from each probe.

Intersection with ME BC00 (“ $\Delta \cap BC00_{ME}$ ”). The ASes from this category were detected as UTRS participants during RAE, but during ME, we did not receive replies to *B-IP* and *C-IP* from them. As we noted in Section 3.2, in ideal conditions, we should not get such results because we had found the corresponding pingable IP by getting a reply to *A-IP*. Therefore, we assume that these cases are due to the flaky nature of Internet measurements.

Summing up all these considerations, the final set of UTRS participants for a measurement obtained using both experiments data is the union of ME UTRS Participants, RAE UTRS Participants not in ME Checked ASes, False Negatives for our Main Experiment (*Intersection with ME BC0x* (“ $\Delta \cap BC0x_{ME}$ ”)) and RAE UTRS Participants that were filtered out during ME because of the absence of ping replies to both our IPs (*Intersection with ME BC00* (“ $\Delta \cap BC00_{ME}$ ”)). The row “Final UTRS Participants” in Table 2 shows the corresponding numbers for each measurement. The final set of highly likely UTRS participants is obtained by taking the intersection of all measurements. Thus, we claim that we are able to detect 562 active UTRS participants with a very high probability. In the following section, we provide the characterization of ASes from this set.

5 UTRS MEMBERS CHARACTERIZATION

In this section, we provide the characterization of the identified UTRS participants. To achieve this goal, we use publicly available external data.

Figure 4 plots the customer cone size as well as the AS rank of the UTRS participants versus non-participants. To build this figure, we enriched our data with the information from the CAIDA’s ASRank [9] dataset. As can be seen, the detected UTRS participants tend to have a higher rank than non-participating ones. Similarly, when looking at the cone size, for both metrics (number of addresses and number of ASes), the detected UTRS participants have proportionally more and bigger customers than those not participating. This signals the willingness of big and medium-sized ASes to join the UTRS.




Fig. 4. AS rank and customer cone size eCDFs

We also classified UTRS participants using ASdb [53]. ASdb provides an updated North American Industry Classification System (NAICSlite) and classifies, according to this system, the registered organizations associated with ASNs. Each ASN in the dataset is associated with one or several Categories (Layer-1) and Sub-categories (Layer-2). ASdb incorporates the data from PeeringDB [34] and from CAIDA AS Classification [10] and supersedes them in terms of coverage and the quality of classification. We downloaded the latest¹ available ASdb dataset [53] and enriched our data using this information. Table 3 contains the obtained results classifying the UTRS participants according to Category 1 - Layer 1 and Category 1 - Layer 2.

¹The dataset snapshotted in May 2022.

Table 3. UTRS participants classification using ASdb (Total: 196)

| Category & Sub-Category | # | Category & Sub-Category | # |
|---|------------|---|-----------|
| Computer and Information Technology | 138 | Education and Research | 19 |
| <i>Hosting and Cloud Provider</i> | 14 | <i>Colleges, Universities, and Professional Schools</i> | 4 |
| <i>Internet Service Provider (ISP)</i> | 100 | <i>Education Software</i> | 2 |
| <i>Software Development</i> | 4 | <i>Other</i> | 1 |
| <i>No sub-category</i> | 20 | <i>No sub-category</i> | 12 |
| Retail Stores, Wholesale, and E-commerce Sites | 4 | Media, Publishing, and Broadcasting | 2 |
| <i>Clothing, Fashion, Luggage</i> | 1 | <i>Online Informational Content</i> | 1 |
| <i>Other</i> | 2 | <i>Radio and Television Providers</i> | 1 |
| <i>No sub-category</i> | 1 | <i>No sub-category</i> | 1 |
| Community Groups and Nonprofits | 4 | Manufacturing | 6 |
| <i>Human Rights and Social Advocacy</i> | 1 | <i>Clothing and Textiles</i> | 4 |
| <i>No sub-category</i> | 3 | <i>No sub-category</i> | 2 |
| Health Care Services | 5 | Service | 5 |
| <i>Hospitals and Medical Centers</i> | 3 | <i>Buildings, Repair, Maintenance</i> | 1 |
| <i>No sub-category</i> | 2 | <i>Law, Business, and Consulting Services</i> | 3 |
| Travel and Accommodation | 2 | <i>No sub-category</i> | 1 |
| <i>Other</i> | 1 | Construction and Real Estate | 1 |
| <i>No sub-category</i> | 1 | <i>Civil Engineering Construction</i> | 1 |
| Other | 8 | Finance and Insurance | 1 |
| <i>No sub-category</i> | 8 | <i>Investment, Portfolio Management, Pensions and Funds</i> | 1 |

Table 3 shows the majority of the detected UTRS participants belong to the *Computer and Information Technology* category. Almost half of all ASes present in ASdb are from *Internet Service Provider (ISP)* and *Hosting and Cloud Provider* sub-categories. These AS types directly benefit from participating in UTRS because this allows them to reduce the size of DDoS attacks on their networks. Nevertheless, we must note that ASdb did not classify all ASes in our dataset: only 196 out of 562 UTRS participants are categorized.

Using the CAIDA's ASRank [9] dataset obtained during the last week of our measurements, we also built a map of the detected UTRS participants. Figure 5 shows the location of these ASes. As you can see from the figure, three main regions, namely, Western Europe, the USA, and Brazil, host most of the detected UTRS participants. While it is not surprising to see the former two because these are the most developed regions in the world, Brazil is an unexpected member. At the same time, a recent study shows [30] that the adoption of the anti-DDoS security best practices (namely, source address validation) in Brazilian ISPs is significantly faster than in the rest of the world.




Fig. 5. UTRS participants map

6 UTRS PARTICIPATION DETERMINANTS

In the previous sections, we identified and characterized UTRS participants, but what makes a network operator join such community-based RTBH? What are the determinants for participating in UTRS? To identify the determinants that drive UTRS participation, we conducted a survey based on Protection Motivation Theory (PMT) [38]. PMT is frequently utilized in behavioral cybersecurity research, with numerous instantiations of the theoretical model used in the literature [6, 11, 19]. PMT is appealing to cybersecurity researchers since it focuses on a threat as well as a prevalent countermeasure to that threat.

There are two primary components of PMT. *Threat appraisal* is concerned with perceived vulnerability and severity. The conditional chance that a DDoS attack will damage if no countermeasures are taken is referred to as *perceived vulnerability*. The *perceived severity* of a prospective attack refers to the possible negative implications to the network produced by DDoS attacks in the context of our study. The premise here is that if a DDoS attack happens as a result of not having joined UTRS, and if this becomes a severe problem, they will consider joining UTRS. *Coping appraisal* in PMT entails: (1) assessing the efficacy of the protective behavior in coping with the threat (response efficacy); (2) believing in one's own ability to manage protective behaviors (self-efficacy); and (3) estimating the costs (including money, time, and energy) and efforts required to perform protective behaviors (perceived response cost). Overall, response efficacy and self-efficacy are projected to strengthen coping appraisal, whereas response cost is expected to decrease it (see Figure 6).




Fig. 6. UTRS-adoption research model. Arrows refer to possible influence relationships.

6.1 Data Collection Instrument

Informed by PMT, we created an instrument consisting of a self-reporting questionnaire, which was completed online by network operators after obtaining written informed consent. This questionnaire has two parts: the first is demographic information and experience, and the second – PMT constructs. We developed the questionnaire based on an extensive review of the literature [20, 32, 41]. Demographic variables included age, professional experience, gender, and region. The experience questionnaire consists of three items (Yes/No) and questions about UTRS participation, usage of alternative RTBH services, and ASes under management.

The items were based on PMT, including perceived vulnerability (3 Items), perceived severity (3 Items), perceived self-efficacy (3 Items), perceived response efficacy (3 Items), response cost (4 Items), perceived reward (3 Items), social norms (3 Items), intention (2 Items). Self-efficacy, response cost, and threat severity were measured with scales adapted from [51]. Effort expectancy, rewards, and perceived severity items were adapted from [14]. Social norm items were adapted from [47]. All items are stated in Appendix B (see Table 8).

The PMT items were measured on a five-point Likert scale ranging from 1 (strongly disagree) to 5 (strongly agree).

Table 4. Correlation results between the self-reported control variables and the endogenous variable

| Control variable | Correlation (ϕ) | p-value |
|------------------|------------------------|---------|
| Gender | 0.13 | 0.45 |
| Experience | 0.18 | 0.74 |
| Region | -0.04 | 0.86 |
| Age | -0.18 | 0.21 |
| Size | -0.01 | 0.79 |

Notes: Cramer's V test computed for multi-categorical variables, point biserial correlation for numerical variables.

6.2 Participants and Sampling

The data was collected between May and June 2022. Participants were network operators who appeared as technical contact points in PeeringDB [34]. Participants were selected using the stratified cluster sampling method. First, all network operators identified as participating in UTRS were included. This allowed us to validate the results of our methodology. From the rest of the network operators, we selected a random sample of 2,000 networks.

Among the total 2,562 network operators that were invited, 2,491 (89.09%) did not participate. Among the remaining 305 participants, 60 (2.15%) were excluded due to missing data on key variables (UTRS participation, managed ASes), yielding a final sample of 245 (8.76%): 59 UTRS participants and 186 – non-participants. Out of the 59 UTRS participants, 58 self-reported themselves to be UTRS members, which serves as validation of the accuracy of our measurement methodology. On the other hand, out of the 186 non-participants, 13 were not sure about their memberships, and the rest just self-reported to be non-members.

6.3 Control variables

To reduce the endogeneity that comes with the existence of potential confounders, we specify a set of control variables accounting for a part of the dependent variable's variance. We studied five control variables (gender, age, region, experience and network size) to investigate their impact on the endogenous variable, UTRS adoption. The first four control variables are directly related to demographic information, while the last control variable relates to the size of the network operated by the respondent. Table 4 presents the results of the correlation tests between the control variables and the endogenous variable. None of the control variables had a statistically significant correlation with the intention to adopt UTRS.

6.4 Factors Driving UTRS Participation

We use Structural Equation Modeling (SEM) to investigate complicated interactions between latent variables [22]. To examine survey data containing behavioral questions, SEM is often utilized. Given our sample population of more than 200 respondents, SEM allows us to build up and verify links between PMT theoretical constructs and their actual indicators. First, we used confirmatory factor analysis to assess the measurement model that included all latent variables. The second step focused on the structural model, which examined the hypothesized links between the components.

Before looking into the results of the model, we computed descriptive statistics of the responses. The percentages of the respondent's choices per item are shown in Figure 7. There is a clear pattern: most of the respondents (>70%) agreed or strongly agreed with all items apart from the ones inversely coded that showed the opposite patterns. This indicates the suitability and reliability of the instrument, as it was already confirmed by Cronbach's alpha ($\alpha > .7$).

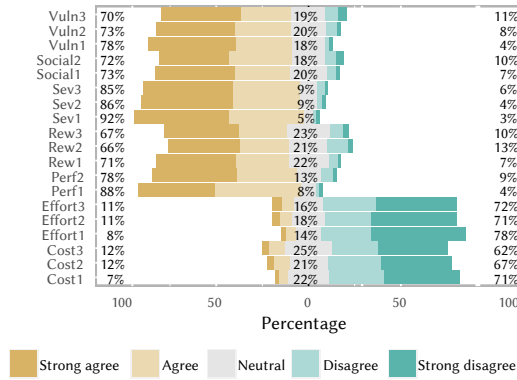


Fig. 7. Survey responses per measurement item

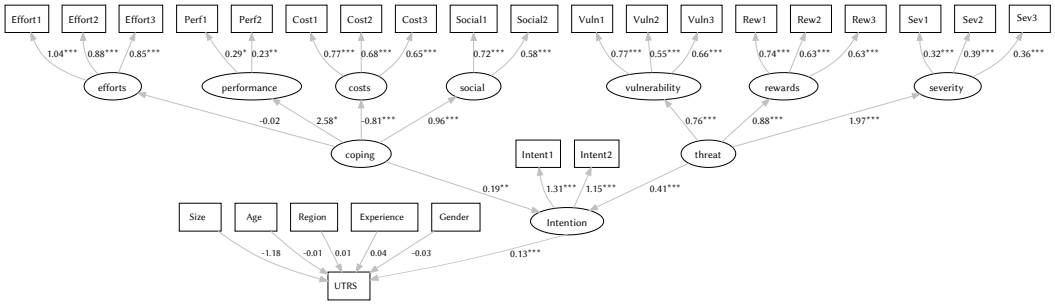


Fig. 8. Fitted SEM path diagram

Table 5. Overview of structural model findings

| | Std. coeff (β) | Std. error | p-value |
|------------------------------------|------------------------|------------|---------|
| <i>Latent variables</i> | | | |
| Severity \rightarrow Threat | 0.89 | 0.04 | 0.00 |
| Vulnerability \rightarrow Threat | 0.60 | 0.05 | 0.00 |
| Rewards \rightarrow Threat | 0.66 | 0.05 | 0.00 |
| Effort \rightarrow Coping | -0.02 | 0.07 | 0.78 |
| Performance \rightarrow Coping | 0.93 | 0.06 | 0.00 |
| Costs \rightarrow Coping | -0.63 | 0.05 | 0.00 |
| Social \rightarrow Coping | 0.69 | 0.05 | 0.00 |
| Threat \rightarrow Intention | 0.37 | 0.06 | 0.00 |
| Coping \rightarrow Intention | 0.17 | 0.06 | 0.00 |
| Intention \rightarrow UTRS | 0.31 | 0.06 | 0.00 |
| <i>Control variables</i> | | | |
| Gender \rightarrow UTRS | -0.02 | 0.06 | 0.68 |
| Experience \rightarrow UTRS | 0.07 | 0.06 | 0.25 |
| Age \rightarrow UTRS | -0.05 | 0.06 | 0.39 |
| Region \rightarrow UTRS | -0.04 | 0.06 | 0.47 |
| Size \rightarrow UTRS | -0.04 | 0.06 | 0.51 |

Next, we look at the standardized coefficients (β) of the fitted structural model. The overview of the findings is provided in Table 5 (see more details in Figure 8). None of the control variables has a

statistically significant impact on the adoption of UTRS. For protection-motivated factors (latent variables), the threat appraisal components, threat severity ($\beta = .89, p < .001$) and vulnerability ($\beta = .60, p < .001$), had positive associations with threat appraisal. Taken together, the results suggest that the threat of DDoS attacks evokes the need for protective measures. Moreover, the rewards of implementing these measures are positively associated with diminishing the threat ($\beta = .66, p < .001$). In turn, the threat appraisal construct ($\beta = .37, p < .05$) is proved to have a positive association with the intention to join the RTBH community. A possible explanation is that network operators see themselves as protectors of network assets and, hence, the threat of DDoS attacks evokes in them the fear to adopt potential countermeasures. Their awareness of security threats and knowledge of the relevant severe consequences are stronger drivers of the adoption of UTRS.

As for coping appraisals of the protection motivation model, the effort to join UTRS ($\beta = -.002, p > .1$) had no statistically significant effect on the intention to participate in UTRS. Response cost ($\beta = -.63, p < .01$) was negatively associated with security protection motivation. On the other hand, social influence ($\beta = -.002, p > .1$) showed a positive effect on the intention to participate.

6.5 Survey Takeaways

Our survey served a twofold purpose: it (i) validated our measurement methodology as most of the respondents voluntarily stated whether they were part of the UTRS, and (ii) helped us gain a deeper understanding of the motivations behind participation in the UTRS. Moreover, the analysis of the survey responses has provided valuable insights into our research question: *What motivates network operators to participate in the UTRS, and what factors might discourage their involvement?*

As stated by some respondents, the UTRS is not a silver bullet to fight DDoS attacks, and as such, it is not a countermeasure that fits every single network operator. In this section, we present the key takeaways from the survey:

- Network operators are more likely to join UTRS if they believe DDoS attacks are a serious threat to their networks. A staggering 78% of these surveyed operators acknowledged the looming threat of DDoS attacks, while an even more resounding 92% concurred that the potential disruptions caused by these attacks were not to be taken lightly.
- Network operators who believe that UTRS can help them protect their networks are less likely to perceive DDoS attacks as a major threat. 88% of respondents voiced their belief that the disruptive force of DDoS attacks could indeed be mitigated through the protective services rendered by UTRS.
- The time and cost required to participate in UTRS can discourage some network operators from joining. A discerning 7% of respondents expressed concern over the potential financial strain imposed by UTRS operations, while 12% harbored reservations about the perceived network sluggishness associated with this security measure.
- Social pressure from peers can encourage network operators to participate in UTRS. A substantial 73% of the surveyed operators found themselves in a network of influence, knowing other operators already were using UTRS. Remarkably, a mere 10% stood as outliers, not perceiving RTBH as the prevailing norm against DDoS attacks.

7 DISCUSSION & LIMITATIONS

UTRS adoption. Our results show that at least 562 network operators from around the globe actively participate in UTRS. Given the collaborative nature of UTRS and the intrinsic network effects that come with it, UTRS would become even more effective if large or even Tier 1 network operators would participate. However, our survey also showed that a major concern for participating is the negative externalities of BGP blackholing, i.e., the impact on legitimate traffic.

On the other hand, the survey showed us that network operators' intentions to participate in UTRS are significantly influenced by perceived vulnerability, perceived severity, response efficacy, potential intrinsic and extrinsic rewards, social norms, and response cost. Only response cost had a substantial negative impact on the intention to participate in UTRS. Because our response cost component is primarily concerned with effort and time required rather than monetary cost, and our coping component captures avoidance behaviors, this may be explained by positing that network operators who estimate a high commitment of time and effort may simply avoid further consideration of UTRS, no matter how desirable it may be in other ways.

Implications and Applications. The measurement methodology and results have practical implications for both researchers and network operators. Researchers can leverage the insights gained from this study to understand the participation patterns in UTRS and enhance their understanding of how network operators utilize RTBH services for DDoS mitigation. Furthermore, the research methodology can be extended to measure and analyze other RTBH services or similar security mechanisms in the future.

For network operators, the results provide valuable information about UTRS support and participation rates among ASes. This knowledge can help network operators make informed decisions regarding UTRS adoption. The characteristics of the participating ASes can significantly impact the efficacy of blackholing. For instance, the presence of Tier 1 networks in the UTRS potentially enhances its effectiveness.

Our findings highlight the significance of allowing network operators to contribute to collaborative RTBH communities. Our research reveals that the operator's impression of how difficult it will be to participate (our response cost construct) has a significant effect on their decision to join UTRS. As the operator's perceived severity of the DDoS threat grows, so does their intention to participate and their proclivity to employ some coping mechanism (such as avoiding the issue). Organizations will need to give both training and assistance to their network operators in order for them to participate effectively in RTBH communities. **Limitations.** One of the trade-offs

of this work is our choice to select only 10 prefixes for each AS to find a pingable IP. According to our measurements, more than 90% of all ASes from the ASRank [9] announce less or equal to this number of prefixes. According to CIDR-Report [2], the mean number of prefixes per AS is 13, with over 26,000 ASes announcing only one prefix. Therefore, we consider this threshold as a reasonable choice that lowers the impact on both the measured and measurement networks. With this threshold, we achieve 93.9% coverage of all Active ASNs.

We select only one pingable IP address per AS. BGP takes IP's prefix into account when selecting a route to this IP. Thus, two packets to two different prefixes of the same AS may flow to their destinations through different paths. However, for this work, we assume that ASes are under a single administrative control and have the same UTRS policy for all its prefixes. Moreover, finding IPs in every prefix creates prohibitive overheads for our work.

Another limitation is that the list of UTRS members is not public. Such a list would provide the ground truth and allow us to verify our findings. Regrettably, Team Cymru informed us that they are unable to provide the list or even confirm the accuracy of ours. However, the results of our survey prove that our methodology is effective in identifying UTRS participants.

In some cases where multiple paths are available to a tested ASN, packets could take either a blocking or a non-blocking path, for example, when using ECMP (Equal-Cost Multipath) [23]. This is not under our control, as this depends on the routing policies and preferences of each AS. However, our analysis should converge, given the large number of traces that we consider in this work. Still, we consider doing an experiment using the Paris traceroute tool [5] in order to prove

our assumptions. However, in this case, we will need to adjust our system setup because, currently, the Paris traceroute tool does not support multiple network interfaces.

8 RELATED WORK

There is a vast body of literature designing and evaluating DDoS countermeasures. In order to mitigate DDoS attacks, multiple technical solutions have been developed with varying degrees of efficiency, hardware or on-premises solutions like F5, Arbor devices installed in the network path that detect and mitigate attacks. These are devices that can address both volumetric and application attacks. However, if the attack volume is above the capacity of the network, or the device itself, the attack cannot be mitigated, causing disruption that possibly extends beyond the attack target [42]. Other solutions focus on applications like WAFs (Web Application Firewall), for example, Akamai or Cloudflare. These solutions hide the actual IP address of the victim website behind their reverse proxy and implement proprietary technologies to filter attacks [39]. Attackers can bypass this type of solution in certain situations due to misconfigurations [25]. Another type of DDoS mitigation involves BGP routing the victim networks, either temporarily during the attack or permanently, through a service provider that offers DDoS protection services, for example, Akamai, Arbor, Cloudflare [43]. [49] proposes a DXP (DDoS Information Exchange Point) that encourages collaboration between networks (IXPs) in detecting and mitigating DDoS attacks closer to the source, as such detection requires visibility at multiple locations. An ISP may also use a hidden Distributed Reflection DoS (DRDoS) honeypot [27] to filter out the unwanted traffic in their networks [52]. In some cases, the costs of these solutions are higher than the benefits, and the victim's ISP may choose to block all traffic to the victim. This is achieved through a mechanism called blackholing or null routing, usually implemented through a BGP RTBH [3, 18]. A study on the efficacy of BGP RTBH at an Internet Exchange is proposed by [33]. However, this is a passive measurement limited to the Internet Exchange providing the passive data.

9 CONCLUSION

In this paper, we present the first Internet-wide investigation of participation in UTRS as a popular DDoS mitigation technique based on RTBH. We designed a methodology for inferring UTRS participation based on active measurements, allowing any UTRS participant to identify who are the other participants. For instance, our methodology can be adopted by Team Cymru in order to promote active UTRS participants.

Our analysis shows that at least 562 networks worldwide actively participate in this RTBH community service to protect their customers and peers. UTRS participants' characteristics show that both large and small networks can benefit from this service. This heterogeneity of participants is also proved by the different sectors in which the participants operate, ranging from education to construction and real estate.

The survey, responded by 245 network operators, shed light on the determinants driving the participation in UTRS. Informed by a theoretical PMT model, our results show that the intention to participate in UTRS is significantly affected by the threat appraisal. Those operators who see themselves as more vulnerable to DDoS attacks are more prone to participate in UTRS. Similarly, operators also showed a coping behavior influencing the decision to participate in UTRS, i.e., operators who foresaw high costs and performance issues for adopting UTRS were more reluctant to join UTRS. Finally, social norms also have a significant impact on the intention to participate.

Acknowledgements. We extend our sincere gratitude to all survey participants. This work is supported by the Dutch Research Council (NWO) under the RAPID project (Grant No. CS.007).

REFERENCES

- [1] Accessed on 03.04.2021. RIPE Atlas. <https://atlas.ripe.net/>
- [2] Accessed on 29.09.2023. CIDR Report. <https://www.cidr-report.org>
- [3] Radu Anghel, Swaathi Vetrivel, Elsa Turcios Rodriguez, Kaichi Sameshima, Daisuke Makita, Katsunari Yoshioka, Carlos H. Gañán, and Yury Zhauniarovich. 2023. Peering into the Darkness: The Use of UTRS in Combating DDoS Attacks. In *European Symposium on Research in Computer Security (ESORICS)*. 23–41.
- [4] Jose Arroyo-Barrigüete, Ricardo Ernst, Jose López-Sánchez, and Alejandro Orero-Giménez. 2010. On the identification of critical mass in Internet-based services subject to network effects. *The Service Industries Journal* 30, 5 (2010), 643–654.
- [5] Brice Augustin, Timur Friedman, and Renata Teixeira. 2007. Multipath tracing with Paris traceroute. In *2007 Workshop on End-to-End Monitoring Techniques and Services*. 1–8.
- [6] Salvatore Aurigemma and Thomas Mattson. 2018. Exploring the effect of uncertainty avoidance on taking voluntary protective security actions. *Computers & Security* 73 (2018), 219–234.
- [7] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The menlo report. *IEEE Security & Privacy* 10, 2 (2012), 71–75.
- [8] Shehar Bano, Philipp Richter, Mobin Javed, Srikanth Sundaresan, Zakir Durumeric, Steven J Murdoch, Richard Mortier, and Vern Paxson. 2018. Scanning the internet for liveness. *ACM SIGCOMM Computer Communication Review* 48, 2 (2018), 2–9.
- [9] CAIDA. [n. d.]. ASRank. Retrieved 05.05.2022 from <https://asrank.caida.org/>
- [10] CAIDA. 2015. AS Classification. Retrieved 25.04.2022 from <https://www.caida.org/data/as-classification/>
- [11] Tim Chenoweth, Robert Minch, and Tom Gattiker. 2009. Application of protection motivation theory to adoption of protective technologies. In *2009 42nd Hawaii International Conference on System Sciences*. IEEE, 1–10.
- [12] Lee J Cronbach and Richard J Shavelson. 2004. My current thoughts on coefficient alpha and successor procedures. *Educational and psychological measurement* 64, 3 (2004), 391–418.
- [13] Christoph Dietzel, Matthias Wichtlhuber, Georgios Smaragdakis, and Anja Feldmann. 2018. Stellar: Network Attack Mitigation Using Advanced Blackholing. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (Heraklion, Greece) (CoNEXT '18)*. Association for Computing Machinery, New York, NY, USA, 152–164. <https://doi.org/10.1145/3281411.3281413>
- [14] Tamara Dinev and Paul Hart. 2004. Internet privacy concerns and their antecedents—measurement validity and a regression model. *Behaviour & Information Technology* 23, 6 (2004), 413–422.
- [15] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. 542–553.
- [16] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *22nd USENIX Security Symposium (USENIX Security 13)*. 605–620. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>
- [17] Xun Fan and John Heidemann. 2010. Selecting Representative IP Addresses for Internet Topology Studies. In *ACM SIGCOMM Conference on Internet Measurement (IMC '10)*. 411–423.
- [18] Vasileios Giotas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. 2017. Inferring BGP Blackholing Activity in the Internet. In *Proceedings of the 2017 Internet Measurement Conference (London, United Kingdom) (IMC '17)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3131365.3131379>
- [19] Anil Gurung, Xin Luo, and Qinyu Liao. 2009. Consumer motivations in taking action against spyware: an empirical investigation. *Information Management & Computer Security* (2009).
- [20] Steffi Haag, Mikko Siponen, and Fufan Liu. 2021. Protection motivation theory in information systems security research: A review of the past and a road map for the future. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems* 52, 2 (2021), 25–67.
- [21] Arman Noroozian Hadi Asghari. 2014. PyPI - pyasn. Retrieved 26.04.2022 from <https://pypi.org/project/pyasn/>
- [22] R Lance Holbert and Michael T Stephenson. 2003. The importance of indirect effects in media effects research: Testing for mediation in structural equation modeling. *Journal of broadcasting & electronic media* 47, 4 (2003), 556–572.
- [23] C Hopps. 2000. RFC992: Analysis of an equal-cost multi-path algorithm.
- [24] Li-tze Hu and Peter M Bentler. 1999. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural equation modeling: a multidisciplinary journal* 6, 1 (1999), 1–55.
- [25] Mattijs Jonker and Anna Sperotto. 2017. Measuring exposure in DDoS protection services. In *2017 13th International Conference on Network and Service Management (CNSM)*. 1–9.
- [26] Daniel Kopp, Christoph Dietzel, and Oliver Hohlfeld. 2021. DDoS never dies? An IXP perspective on DDoS amplification attacks. In *International Conference on Passive and Active Network Measurement*. Springer, 284–301.
- [27] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishioe, Takashi Koide, katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proceedings of the 18th*

- International Symposium Research in Attacks, Intrusions, and Defenses*. 615–636.
- [28] John Kristoff. 2015. *An Internet-wide BGP RTBH Service*. Technical Report. https://www.iab.org/wp-content/IAB-uploads/2015/04/CARIS_2015_submission_20.pdf
 - [29] Linux Kernel. [n. d.]. Linux Kernel Networking Documentation - sysctl. <https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt>
 - [30] Qasim Lone, Alisa Frik, Matthew Luckie, Maciej Korczynski, Michel van Eeten, and Carlos Gañán. 2022. 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)*.
 - [31] Pedro R. Marques, Jared Mauch, Nischal Sheth, Barry Greene, Robert Raszuk, and Danny R. McPherson. 2009. Dissemination of Flow Specification Rules. RFC 5575. <https://doi.org/10.17487/RFC5575>
 - [32] Philip Menard, Gregory J Bott, and Robert E Crossler. 2017. User motivations in protecting information security: Protection motivation theory versus self-determination theory. *Journal of Management Information Systems* 34, 4 (2017), 1203–1230.
 - [33] 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 *Proceedings of the Internet Measurement Conference*. 435–448.
 - [34] PeeringDB. 2004. The Interconnection Database. Retrieved 25.04.2022 from <https://www.peeringdb.com/>
 - [35] Robert A Peterson. 1994. A meta-analysis of Cronbach's coefficient alpha. *Journal of consumer research* 21, 2 (1994), 381–391.
 - [36] Riccardo Ravaioli, Guillaume Urvoy-Keller, and Chadi Barakat. 2015. Characterizing ICMP rate limitation on routers. In *2015 IEEE International Conference on Communications (ICC)*. IEEE, 6043–6049.
 - [37] Yakov Rekhter, Susan Hares, and Tony Li. 2006. A Border Gateway Protocol 4 (BGP-4). RFC 4271. <https://doi.org/10.17487/RFC4271>
 - [38] Ronald W Rogers and Steven Prentice-Dunn. 1997. Protection motivation theory. (1997).
 - [39] Daniel Shugrue. 2017. Fighting application threats with cloud-based WAFs. *Network Security* 2017, 6 (2017), 5–8.
 - [40] Team Cymru. [n. d.]. Unwanted Traffic Removal Service. Retrieved 26.02.2023 from <https://www.team-cymru.com/ddos-mitigation-services>
 - [41] Hsin-yi Sandy Tsai, Mengtian Jiang, Saleem Alhabash, Robert LaRose, Nora Rifon, and Shelia Cotten. 2016. Understanding online safety behaviors: A protection motivation theory perspective. *Computers & Security* 59 (2016), 138–150.
 - [42] Christos Tselios, George Tsolis, and Manos Athanatos. 2019. A Comprehensive Technical Survey of Contemporary Cybersecurity Products and Solutions. In *Computer Security*. Springer, 3–18.
 - [43] Tony Miu Tung, Chenxu Wang, and Jinhe Wang. 2018. Understanding the behaviors of BGP-based DDoS protection services. In *International Conference on Network and System Security*. 463–473.
 - [44] Doughan Turk. 2004. Configuring BGP to Block Denial-of-Service Attacks. RFC 3882. <https://doi.org/10.17487/RFC3882>
 - [45] University of Oregon. [n. d.]. Route Views Project. Retrieved 25.04.2022 from <http://www.routeviews.org/>
 - [46] Jeroen Van Der Ham. 2017. Ethics and Internet measurements. In *2017 IEEE Security and Privacy Workshops (SPW)*. IEEE, 247–251.
 - [47] Viswanath Venkatesh, Tracy A Sykes, and Xiaojun Zhang. 2011. 'Just what the doctor ordered': a revised UTAUT for EMR system adoption and use by doctors. In *2011 44th Hawaii international conference on system sciences*. IEEE, 1–10.
 - [48] Curtis Villamizar, Ravi Chandra, and Dr. Ramesh Govindan. 1998. BGP Route Flap Damping. RFC 2439. <https://doi.org/10.17487/RFC2439>
 - [49] Daniel Wagner, Daniel Kopp, Matthias Wichtlhuber, Christoph Dietzel, Oliver Hohlfeld, Georgios Smaragdakis, and Anja Feldmann. 2021. United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. ACM. <https://doi.org/10.1145/3460120.3485385>
 - [50] Matthias Wichtlhuber, Christoph Dietzel, and Thomas King. 2018. Computer-implemented procedure to defend against or mitigate DDoS attacks on IT infrastructures. <https://patents.google.com/patent/DE102018130588B4/en>
 - [51] Michael Workman, William H Bommer, and Detmar Straub. 2008. Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in human behavior* 24, 6 (2008), 2799–2816.
 - [52] Yury Zhauniarovich and Priyanka Dodia. 2019. Sorting the Garbage: Filtering Out DRDoS Amplification Traffic in ISP Networks. In *Proceedings of the IEEE Conference on Network Softwarization (NetSoft)*. 142–150. <https://doi.org/10.1109/NETSOFT.2019.8806653>
 - [53] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. 2021. ASdb: a system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference*. 703–719.

APPENDICES

A ETHICS

We had a detailed discussion with the university's IRB and received clearance to conduct this study. We conducted our own scans since there is no existing public dataset that reveals UTRS membership. Our study followed all the active monitoring guidelines for ethical network measurement research [46], including creating a web page running at the IP address of the scanner, communication with Internet response teams, and providing an opt-out option for network operators. We want to stress that our measurement did not create any burden on the UTRS participants as once they joined UTRS, dropping the traffic becomes an automated process. We also randomly distributed our queries across the IPv4 address space, so the scanner to detect pingable IP addresses does not consistently query the same AS before moving on to the next one. Furthermore, in line with the Menlo report [7], we considered that the marginal impacts of these measurements are outweighed by the beneficence of measuring the effectiveness of UTRS and potentially attract more participants..

The institutional ethics committee also cleared our survey instrument with was thoroughly evaluated and approved by them. Survey participants were briefed about the data collection process through the consent form at the beginning of the study. Additionally, at the end of the study, all participants were debriefed to clarify that their data will not be shared with anyone other than the researchers conducting this study.

B SURVEY: DEMOGRAPHICS

Table 6. Survey participant characteristics

| Variable | UTRS participation | | p-value ² |
|-------------------|--------------------------|--------------------------|----------------------|
| | No, N = 176 ¹ | Yes, N = 69 ¹ | |
| Gender | | | 0.46 |
| Female | 1 (0.6%) | 0 (0%) | |
| Male | 164 (93%) | 65 (94%) | |
| Others | 5 (2.8%) | 0 (0%) | |
| Prefer not to say | 6 (3.4%) | 4 (5.8%) | |
| Age | | | 0.99 |
| 18 - 24 | 2 (1.1%) | 1 (1.4%) | |
| 25 - 34 | 14 (8.0%) | 6 (8.7%) | |
| 35 - 44 | 79 (45%) | 32 (46%) | |
| 45 - 54 | 52 (30%) | 20 (29%) | |
| 55 - 64 | 28 (16%) | 10 (14%) | |
| 65 - 75 | 1 (0.6%) | 0 (0%) | |
| Experience | | | 0.19 |
| 6 months - 1 year | 7 (4.0%) | 1 (1.4%) | |
| 1 year - 3 years | 37 (21%) | 15 (22%) | |
| 3 years - 5 years | 71 (40%) | 37 (54%) | |
| Over 5 years | 61 (35%) | 16 (23%) | |
| Region | | | 0.89 |
| Africa | 3 (1.7%) | 2 (2.9%) | |
| Asia | 9 (5.1%) | 3 (4.3%) | |
| Europe | 78 (44%) | 34 (49%) | |
| Latin America | 14 (8.0%) | 4 (5.8%) | |
| North America | 72 (41%) | 26 (38%) | |

¹ Median (IQR) and Frequency (%)

² Fisher's exact test

C ASSESSMENT OF THE MEASUREMENT MODEL

We used reliability, convergent validity, and discriminant validity metric to evaluate the measurement model. First, we evaluated the reliability by checking the reliability of the components by using Cronbach's alpha [12]. The reliability of entire structures exceeds .70, which is a usually utilized threshold [35]. The numerical value is between .822 and .886, both of which are above .70, as illustrated in Table 7. This outcome indicates sufficient reliability. Secondly, to evaluate convergence validity, we looked at Extracted Average Variance (AVE), which should exceed the .50 threshold. It can be seen from Table 7 that the AVE value exceeds .50 for all items. Therefore, our model has a good convergence validity. Finally, the validity of the discrimination if the composite reliability is greater than .7, which is the case for all constructs.

Table 7. Measurement model reliability of latent variables.

| | Latent variables | | | | | | |
|-----|------------------|-------|-------|-------|-------|-------|-------|
| | SEV | VUL | EFF | PER | REW | COS | SOC |
| CA | 0.822 | 0.839 | 0.840 | 0.853 | 0.846 | 0.886 | 0.879 |
| CR | 0.843 | 0.845 | 0.829 | 0.880 | 0.849 | 0.885 | 0.887 |
| AVE | 0.649 | 0.646 | 0.618 | 0.790 | 0.653 | 0.720 | 0.798 |

Notes: CR = Composite Reliability, AVE = Average Variance Extracted, CA = Cronbach's Alpha, SEV = Threat Severity, VUL = Threat Vulnerability, EFF = Self-efficacy, PER = Response efficacy, REW= Intrinsic and extrinsic rewards, COS = Response Cost, SOC= Social influence

Next, we tested the whole model. First, we evaluated the research model's global goodness-of-fit. A good model should have a root mean square error of approximation (RMSEA) <.06, a comparative fit index (CFI) >.95, Tucker Lewis Index (TLI) >.95, and a standardized root mean square residual (SRMR) <.08 [24]. In our study, the results showed an excellent approximate fit for the research model: RMSEA = .01 (90% C.I. = .000, .033), CFI = .999, TLI = .998, and SRMR = .072. The chi-square statistics were significant ($\chi^2(199) = 305.99, p < .001$).

D NETWORK OPERATORS SURVEY INSTRUMENT

Table 8. The measurements and items

| Latent | Items |
|----------------------|---|
| Severity | |
| Sev1 | DDoS attacks are a serious security threat |
| Sev2 | DDoS attacks can cause serious complications including service downgrade and downtime. |
| Sev3 | DDoS attacks can cause reputation damage. |
| Vulnerability | |
| Vul1 | My organization could get a DDoS attack. |
| Vul2 | My network will get disrupted if it gets a DDoS attack. |
| Vul3 | The frequency of DDoS attacks has increased over the last few years. |
| Self-efficacy | |
| Effort1 | It is easy to mitigate DDoS attacks through remote blackholing. |
| Effort2 | I would join a remote blackholing community regardless of its cost. |
| Effort3 | To mitigate a DDoS attack, I would drop traffic despite the possibility of the side effects (e.g., dropping legitimate traffic). |
| Performance | |
| Perf1 | DDoS disruptions can be mitigated by joining a collaborative remote blackholing service like the Unwanted Traffic Removal Service (UTRS). |
| Perf2 | Collaborative remote blackholing is one of the best solutions for counteracting problems caused by DDoS. |
| Rewards | |
| Rew1 | I will help others if I join a collaborative remote blackholing service like the Unwanted Traffic Removal Service (UTRS). |
| Rew2 | Collaborative remote blackholing is more convenient to take other countermeasures to prevent DDoS attacks. |
| Rew3 | It will save me time/money if I join a collaborative remote blackholing service like UTRS. |
| Costs | |
| Costs1 | Community-based remote blackholing is expensive/costly to configure and operate. |
| Costs2 | Collaborative blackholing requires updating the configuration all the time. |
| Costs3 | Collaborative remote blackholing can slow down your network |
| Social | |
| Social1 | I know several network operators that use a collaborative remote blackholing service like UTRS. |
| Social2 | Community-based blackholing is the norm among network operators to mitigate DDoS. |
| Intention | |
| Intent1 | I intend to use a community-based remote blackholing service in the near future. |
| Intent2 | I expect that community-based remote blackholing continues to be of use in the future. |

Received August 2023; revised January 2024; accepted January 2024