Glossy Mirrors: On the Role of Open Resolvers in Reflection and Amplification DDoS Attacks

Ramin Yazdani*, Max Resing‡, Anna Sperotto*
*{r.yazdani, a.sperotto}@utwente.nl, *University of Twente*
‡{max.resing}@netscout.com, *NETSCOUT*

***Abstract*—Open DNS resolvers are infamous contributors to DDoS attacks. Characteristics of open DNS resolvers have been studied in different aspects in the past. However, there is a gap in knowledge on the actual role of open resolvers acting involuntarily as DNS reflectors in DDoS attacks.**

In this paper, we study DNS reflectors in more than half a million DDoS events using a large-scale DDoS telemetry dataset provided by a DDoS protection service provider with a global footprint. Our findings reveal that while the majority (∼79%) of reflectors misused in attacks are open resolvers capable of delivering large DNS responses, the contribution of reflectors with very small response sizes is not negligible either. Additionally, our analyses reveal that the distribution of misused open resolvers is biased toward certain countries and network operators, likely impacted by the IP churn of reflectors, while in terms of network types, there is no outstanding bias visible in an aggregated view. Finally, comparing the pool of misused open resolvers to the pool of all exposed and potentially abusable resolvers, the latter dwarfs the former, suggesting that the firepower of DNS-based DDoS attacks could substantially increase in the future.

I. Introduction

Distributed Denial of Service (DDoS) has been a persistent and ever-growing threat to the availability of networks and services on the Internet [8], [12], [19]. Reflection & Amplification (R&A) attacks [25] are one of the popular DDoS attack types [18] and the DNS protocol is one of the most common attack vectors for this attack type [8], [12], [26], [29].

DNS-based reflection/amplification DDoS attack events typically misuse open DNS resolvers by sending them queries with spoofed source addresses. These resolvers in return send a response (which is typically larger than the query in size) to a victim, and, in orchestration, can exhaust the network's capacity of the victim or its upstream infrastructure.

The pool of exposed open resolvers is significantly larger, by multiple orders of magnitude, than the typical number of reflectors that are misused in attacks [19], [34]. This raises the question if there is any rationale behind the selection of the exploited reflecting infrastructure. However, there is limited information available about reflectors that are actively exploited in attacks and their characteristics.

In this paper, we investigate the misuse of open resolvers in DDoS events in practice to understand patterns and if adversaries possibly apply any selection criteria. Understanding those patterns could help improve current operational practices in fighting against DNS R&A attacks. Additionally, we aim at finding out how efficient adversaries are in selecting the ”right” infrastructure - namely open resolvers reflecting with large response sizes - and to what extent their firepower can grow when changing their pool of open resolvers. By studying hundreds of thousands of DNS R&A DDoS events, we present the following contributions:

- We shed light on the contribution of different types of DNS reflectors to DDoS events and show that while open resolvers properly processing DNS queries are the main contributors to DNS R&A DDoS events, the number of misused resolvers with a failed or anomalous response also have a non-negligible impact.
- We investigate the correlation between the number of misused reflectors and the impact of DDoS events and show that the intensity of DDoS events has no linear dependency on the number of misused reflectors.
- We study the demographics of the misused reflectors and show that actual misuse of reflectors is biased towards certain countries compared to the general exposure of open resolvers in the wild. We also discovered that IP churn (the rate at which the IP addresses of hosts change in a network) plays a role here. This means that countries with a high churn on open resolver IPs are not among the countries with the most abused infrastructure, whereas low IP churn makes it more likely to show up on that list.

The remainder of this paper is structured as follows. In Section II, we provide details about the main datasets used in our work. In Sections III to V, we elaborate on our findings. In Section VI, we provide some operational recommendations and highlight the limitations of our work. We present related work in Section VII. Finally, we conclude our paper in Section VIII.

II. Datasets

We make use of two datasets. The first is a DDoS telemetry dataset (Sec. II-A). The second encompasses a set of weekly scans for open DNS resolvers (Sec. II-B).

A. NETSCOUT Sightline

NETSCOUT is a network and application performance management company that primarily serves Internet service providers, data centers, and enterprises alike requiring comprehensive network observability and analysis capabilities. In recent years, their products have been complemented with bidirectional threat intelligence sharing, which improves the

978-3-903176-66-9 ©2024 IFIP

security features of said products. NETSCOUT's products are present in hundreds of networks spread across multiple continents, providing representative insights into the global threat landscape.

NETSCOUT's Sightline is a DDoS attack detection and mitigation solution positioned in the center of a network, receiving network flow telemetry from the various infrastructures on a network's edge. Their products were complemented with bi-directional anonymized threat intelligence sharing. Aggregations of this telemetry, in particular about ongoing DDoS events (classified as alerts heretofore), such as timestamps, protocols, impact, attack vectors, origin (attacking IPs), and anonymized destination of the traffic, are looped back to NETSCOUT's threat intelligence infrastructure.

NETSCOUT's DDoS telemetry is a proprietary dataset. However, NETSCOUT has a dedicated research team which collaborates with researchers who are working on the analysis of DDoS threats, trends and insights. In the following sections, we elaborate on details of the Sightline DDoS event telemetry dataset that we use in this paper.

1) Demographics of Telemetry Sources: Sightline is deployed and operating in more than 500 Autonomous Systems (ASes) in 95 different nations around the globe. This global network observes a combined average of 500Tbps of peering traffic during the second half of 2023 [21]. The biannual DDoS threat intelligence report also attests to a grand total of over 13.1 million DDoS attacks in 2023 alone. Based on IANA's ASN allocation statistics[1], Sightline's coverage is roughly proportional to the number of allocated ASNs per RIR. Considering that Sightline is primarily designed for service provider networks, we see a strong industry bias towards wired and wireless telecommunication carriers as well as cloud, hosting, and collocation providers.

2) DDoS Attack Telemetry: It is not uncommon for adversaries to take advantage of combining multiple DDoS attack vectors [5], [12]. Sightline is able to attribute multiple attack vectors to the same alert instance. The sensitivity depends on a network- and Sightline-optimized configuration. Due to the high-level aggregation of DDoS events, we filter on alerts with a high confidence of being single vector DNS R&A events, originating from UDP source port 53 (the default DNS port). Additionally, we focus on incoming alerts only - from the perspective of the edge routers of a network.

3) Source Address Aggregation: The DDoS attack telemetry submitted by Sightline (referred to as the *primary* dataset in the remainder of this paper) aggregates source IP addresses involved in attacks for performance reasons, keeping track of roughly the top 200 source IPs per DDoS alert. In this work, we focus on the non-aggregated IPv4 source addresses of the input dataset. Nawrocki et al. [19] reported that 90% of DNS R&A attack events leverage less than 100 amplifiers. Thus, we argue that this dataset with the most-impactful 200 host IP addresses per attack event provides a representative picture of DNS reflectors abused in DDoS attacks. We further evaluate the impact of source address aggregation in Section II-A4.

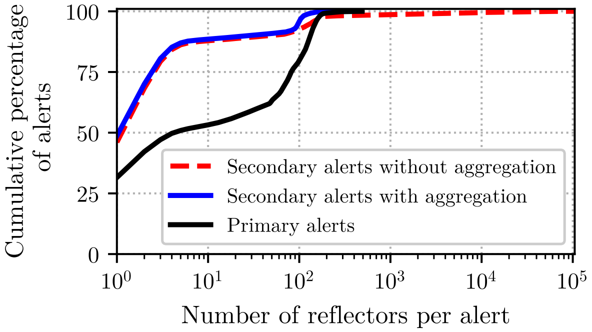

Fig. 1: The cumulative distribution of the number of reflectors per alert. The vast majority of alerts involve less than 200 reflectors.

4) Secondary Insights: In the period December 2023, we have access to a second dataset, namely DDoS attack events with the full list of involved source IPs. Due to externalities, this dataset is limited in scope and time. Additionally, it only covers a subset of sensors for the full set of DDoS attack sources. We use this data to verify that our observations are generalizable for the dataset with the primary DDoS attack event telemetry. In the remainder of this paper, we refer to this dataset as *secondary insights*.

In Fig. 1, we compare the distribution of the number of reflectors per alert for alerts in the *secondary insights* data set in which no source IP aggregation was performed (dashed red line), to the corresponding alerts with source aggregation extracted from the *primary* dataset (solid blue line). We observe that the two datasets show an almost identical distribution for around 85% of the DDoS alerts. This confirms that by relying on the dataset with source address aggregation we only miss a small part of information on the misused reflectors. We also plot the distribution of the number of reflectors per alert for our *primary* dataset spanning over 2023 in Fig. 1. We observe that alerts in the primary dataset have a higher number of reflectors per alert compared to our *secondary insights* dataset.

B. Open DNS Resolver Scans

We further rely on our longitudinal dataset of open DNS resolver discovery scans to explore reflectors that are misused in DDoS attacks. Due to security and ethical considerations, we do not publicize this dataset. However, there are services such as Censys [2] that provide researchers with similar datasets. We collect weekly snapshots of hosts with port 53 open in the IPv4 address space. This list includes traditional open DNS resolvers, transparent forwarders [20], as well as any host that returns a DNS response with an arbitrary DNS response code. In the case of transparent forwarders, it is possible to consider the IP address of the forwarder, or the responding host, as an open resolver. From the scanner point of view, we can consider the probed IP address as an open

[1] https://www.iana.org/numbers/allocations/

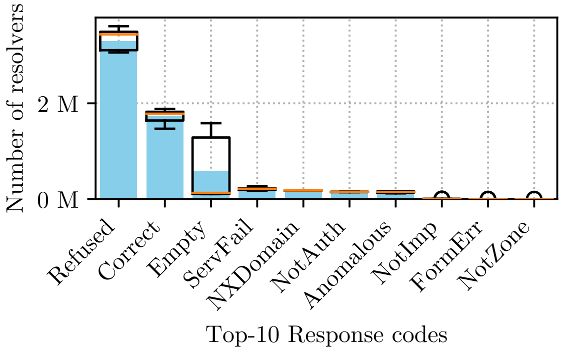

Fig. 2: Top-10 DNS response codes returned by open DNS resolvers over 2023. The boxplots show the statistical variability in our weekly DNS scans.

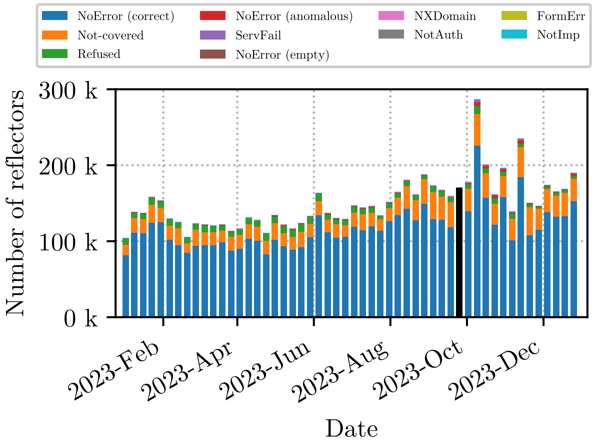

Fig. 3: The number of reflectors seen in R&A DDoS alerts per week over 2023 subdivided based on the response code these resolvers return in the scans. The black bar in September corresponds to a week where we do not have an open resolvers scan snapshot.

resolver. However, from a victim's point of view, the answer comes from the ultimate recursive resolver without being aware of the existence of the forwarder. In this work, we use responding source IPs rather than probed IPs to account for transparent forwarders. This is because the input dataset of attack traces reports hosts that send the response to the victim. This approach comes with a limitation as there is no one-to-one mapping between transparent forwarders and their corresponding recursive resolver.

In Fig. 2, we plot the number of resolvers that return different DNS rcodes in their responses. In this plot, we consider the responding IP addresses, which are not necessarily the same as the probed open resolver. On average, we observe roughly 1.7M distinct responding IP addresses that return a correct/expected answer (the resource record that our authoritative nameserver returns) to our DNS queries, referred to as *open recursive resolvers* in the rest of this paper. From a DDoS attacker's point of view, these resolvers are more appealing since they can be leveraged to amplify the DNS response arbitrarily. On the other hand, roughly 3.3M hosts respond with a `Refused` answer. We argue that the majority of the `Refused` answers correspond to authoritative nameservers. While these resolvers can also be misused to reflect DNS responses, the responses they return are typically small in size and thus do not result in amplification.

In the rest of this paper, we refer to all IPs responding to UDP port 53 probes that are not a recursive resolver as a *non-recursive resolver/reflector*. Some of these hosts may be recursive resolvers that simply block queries issued by the open resolver discovery scans used in this paper. We suspect that this should rarely be the case. However, we leave a thorough investigation of this as a future work.

III. Attack Statistics

As a starting point for our analyses, we draw statistics about the number of DDoS alerts in our input dataset. Only 0.86% of alerts and 1.18% of reflectors in our *primary* dataset are IPv6 alerts and reflectors. Thus, in the rest of this paper, we only focus on IPv4 alerts and reflectors. In total, we see 555,881 incoming DNS R&A DDoS attacks over the year 2023 with a gradually increasing trend in the number of attacks over time.

A. Reflectors Analysis

We cross the source IP addresses in DDoS alerts with the open DNS resolver dataset and plot the number of reflectors seen in DDoS alerts aggregated per week in Fig. 3. We break down source IP addresses into different categories based on the status of the response that they return in the open resolver scans. Note that the DNS queries in the open resolver scan dataset are not the same as the ones used in DDoS attacks. This means that non-recursive resolvers in the open resolver scan dataset can return a different answer in practice, based on the query they receive. For example, an authoritative nameserver would respond to the open resolver scans with a `Refused` answer, while it can successfully resolve the domain name that it is authoritative for.

We observe an increasing trend of the number of reflectors per week which is in line with the increasing number of DDoS alerts per week. Thus, the average number of reflectors per alert is quite stable over time. While the majority of sources in DNS R&A DDoS alerts (∼117.9 k in average per week) correspond to open recursive resolvers that successfully return a correct answer in our scans (i.e., blue bars in Fig. 3), we observe a non-negligible number (∼23.2 k in average per week) of sources that are not covered in the scans (orange bars). We attribute this group of attack sources to be mostly spoofed sources (further investigated in Section III-B). Additionally, a non-negligible number (∼6.7 k in average per week) of sources return a `Refused` response. We argue that the majority of these hosts are authoritative nameservers that would not perform recursion for queries that they are not authoritative for. Considering that our input dataset does

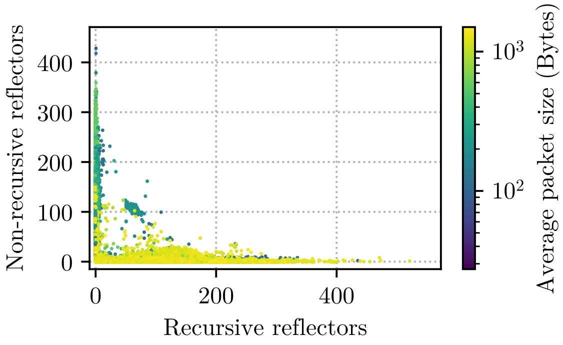

Fig. 4: The scatter plot of the contribution of reflectors with an expected DNS response and all other reflectors to DDoS alerts through 2023 limited to alerts with an average packet size of 1500 Bytes to increase readability. Alerts comprising mostly open recursive resolvers typically have a higher intensity compared to alerts that mainly involve non-recursive reflectors.

not contain DNS query names or individual packet sizes, we cannot infer how these hosts were misused in attacks. Finally, Fig. 3 shows a growth in the misused reflectors that return an anomalous response. These resolvers are also not very appealing in DDoS attacks as the attacker has no control over the expected response size.

Observing the involvement of non-recursive reflectors in DNS R&A DDoS attacks is an unexpected event. Thus, we examined the contribution of these hosts to DDoS attacks by breaking down the number of reflectors per category (recursive and non-recursive) in each individual alert in Fig. 4, color-coded with the average packet size (derived from dividing the average bandwidth by the average packet rate) in an alert. We observe three clear patterns for the composition of attacks. First, a horizontal pattern corresponding to attacks in which mostly reflectors with an expected answer are misused. Arguably, this is the most efficient way to leverage reflectors to bring about DDoS attacks since attackers can achieve a high bandwidth amplification factor using a low number of reflectors. The second pattern is a vertical pattern corresponding to attacks in which mostly reflectors with incorrect/unexpected DNS responses are involved. In terms of bandwidth amplification, these attacks are the least efficient. This observation supports our intuition that those sources, albeit reacting to the DDoS probes, are not effectively contributing to the attack. The third pattern we observe corresponds to alerts that roughly use the same number of reflectors within each category.

We further investigate the residual DNS reflection potential in terms of the number of available reflectors by comparing the number of misused reflectors per alert to the number of all exposed reflectors during that event. We perform this comparison separately for open recursive resolvers (i.e., reflectors that return the expected response), and all other hosts that respond to our scans with an incorrect/unexpected answer. For this purpose, we rely on our *secondary* dataset that reports the full list of reflectors in an event. We observe that the DDoS alerts with the highest number of misused resolvers leverage ∼94 k open resolvers and ∼2.4 k non-recursive resolvers. These numbers are substantially lower than the total number of exposed open resolvers (5.9% of ∼1.6 M) and non-recursive resolvers (0.06% of ∼4.3 M) during the same time period. Considering that the vast majority of DDoS alerts comprise less than 200 reflectors (see Fig. 1) raises a concern that the number of misused reflectors in DNS R&A attacks could potentially grow by multiple orders of magnitude. In terms of attack intensity (i.e., bandwidth and packet rate), there are instead additional elements than the number of misused reflectors in defining the untapped potential of DNS R&A DDoS attacks, such as the upstream network capacity or the presence of rate limiting. This makes a direct estimation of the possible growth of attack intensity not straightforward, but data on the number of available reflectors and current attack intensities suggest that growth is very well possible.

The same DNS reflectors are typically involved in many distinct attacks [13]. We investigate the reuse of reflectors across multiple alerts. We observe that 29% of reflectors are present in only one DDoS alert. This percentage is lower than one reported by Nawrocki et al [19] where 50% of resolvers were involved in a single DDoS event. We argue that this is partially because of the longer time span of our dataset such that there is a higher chance of reuse of reflectors. Alternatively, 99% of reflectors are present in less than 0.04% of alerts in our study. A handful of reflectors comprising large public DNS providers contribute to the long tail of more frequently misused resolvers, by being involved in at most 10.5% of DDoS alerts.

B. *Not-covered Attack Sources*

As we observe in Fig. 3, a non-negligible number of attack sources are not listed in the open resolvers scan dataset. A probable explanation for this artifact is that these sources are randomly spoofed. Randomly spoofed DNS packets (or crafted packets with source port 53) can impersonate typical DNS responses, and therefore their source IPs would be logged as contributors to an attack. Another underlying cause behind the not-covered sources may be that these are resolvers in networks that block the open resolver scans. To evaluate whether this is the reason behind our observation, one needs to run scans from different vantage points, which is out of the scope of this paper. In the following, we investigate a number of features that can indicate the spoofed nature of these attack sources.

First, we look at the temporal persistence of these sources in our dataset. In Fig. 5, we plot the cumulative distribution of the number of weeks in which the not-covered attack sources are observed and compare it to the distribution of the rest of the sources. While open DNS resolvers are known to have a high IP churn [3], [15]–[17], [28], [34], we observe that the sources that are not covered by our scans have an even shorter lifetime in attacks. 60% of open resolvers involved in attacks are seen on more than a single week of alerts, while for not-covered sources this is only 35%.

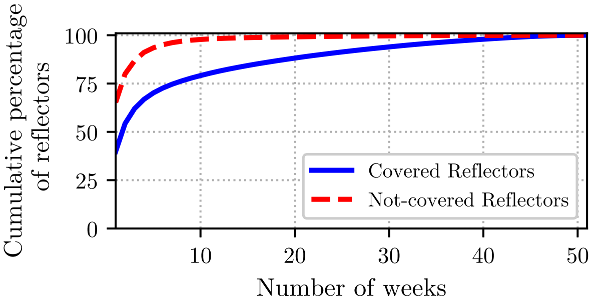

Fig. 5: The distribution of the number of weeks in 2023 in which reflectors are seen for reflectors in alerts covered in our open resolver scans compared to reflectors in alerts that are not covered in our scan data. Not-covered reflectors have a relatively shorter lifetime compared to reflectors covered in DNS scans, suggesting that they might be randomly spoofed sources.

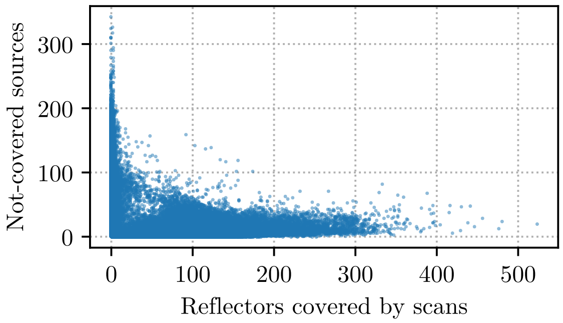

Fig. 6: The scatter plot of the contribution of the number of reflectors covered in our open resolver scans and all other attacker source IPs in DDoS alerts through 2023. We attribute alerts with mostly not-covered sources to randomly spoofed attacks.

Next, we look at the contribution of alerts per category to individual DDoS alerts. In Fig. 6, we plot the number of reflectors per alert that are observed in our DNS scans against the number of not-covered sources in the same alert. We observe two common patterns in this plot. First, a horizontal pattern corresponding to attacks in which mostly reflectors are misused that covered in our scan dataset. The second pattern is a vertical one, corresponding to alerts in which the majority of reflectors are not listed in our scan dataset. The extreme case corresponds to an alert in which none of the 342 reported sources are listed in our scan dataset. We argue that this observation indicates the likely spoofed nature of these sources.

We also observe a number of sources that belong to reserved IPv4 address ranges (e.g., IPs from Private-Use and Loopback ranges), despite the fact that intermediate routers should drop these packets in transit. Additionally, we observe source IP addresses belonging to the UCSD Network Telescope [1] range in our alerts, which we verified as being dark IPs. We consider this an additional confirmation of our hypothesis that these not-covered sources are randomly spoofed source addresses. We exclude these sources from our analysis in the rest of this paper as they do not provide reliable insights about the ecosystem of misused reflectors.

IV. Attack Impact

NETSCOUT's DDoS dataset also reports the intensity of each DDoS event in terms of packet rate (pps) and bandwidth (bps). We use this data to study the correlation between the number of misused DNS reflectors and attack intensity as shown in Fig. 7. Additionally, we plot an upper limit (solid diagonal red line) and a lower limit (dashed diagonal red line) in Fig. 7. The solid diagonal line shows the standard Ethernet MTU (1500 bytes) and the dashed diagonal line shows the theoretical lower limit for a DNS packet (40 bytes).

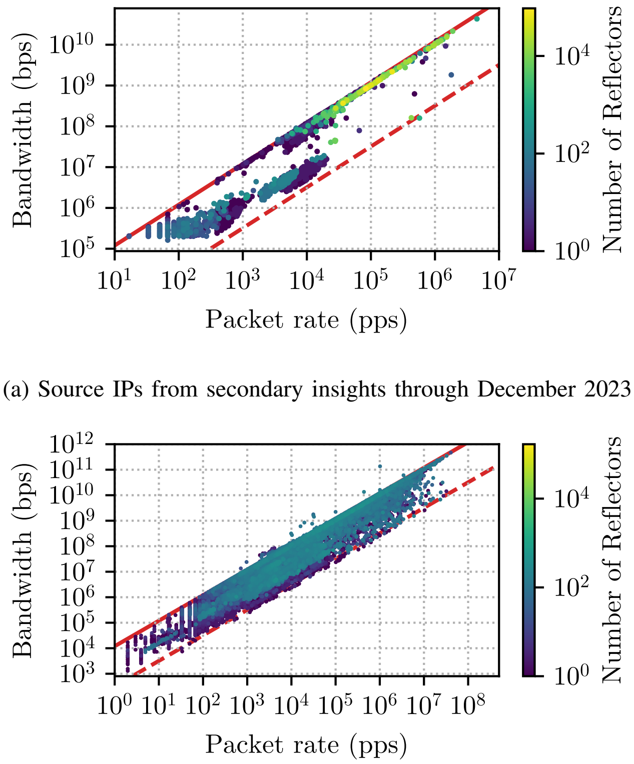

(a) Source IPs from secondary insights through December 2023

(b) Source IPs from primary DDoS telemetry through 2023

Fig. 7: Scatter plot of the number of reflectors per alert compared to the attack traffic and packet rate. The solid diagonal line shows the standard Ethernet MTU (1500 bytes) and the dashed diagonal line shows the theoretical lower limit for a DNS packet (40 bytes).

We start investigating the impact of attacks by examining the *secondary insights* dataset where there is a complete visibility into all reflectors misused in attacks. We plot this in Fig. 7a.

TABLE I: Top-10 countries hosting open DNS resolvers compared to top-10 countries of open resolvers that are misused in attacks from January to November 2023.

	All open resolvers				Misused open resolvers		
Country	Median (%)	Mean (%)	Std (%)	Country	Median (%)	Mean (%)	Std (%)
CN	543.2 k (30.33%)	518.9 k (29.71%)	65.6 k (12.64%)	RU	18.2 k (16.47%)	19.4 k (16.67%)	4.1 k (21.04%)
IN	154.9 k (8.65%)	149.2 k (8.54%)	13.8 k (9.28%)	ID	11.9 k (10.78%)	12.2 k (10.49%)	2.6 k (21.41%)
US	136.0 k (7.59%)	133.0 k (7.62%)	13.3 k (10.00%)	BR	6.4 k (5.82%)	6.6 k (5.65%)	1.0 k (15.87%)
RU	93.6 k (5.22%)	93.0 k (5.33%)	4.4 k (4.70%)	UA	6.1 k (5.50%)	6.5 k (5.57%)	1.4 k (22.13%)
KR	91.4 k (5.11%)	91.6 k (5.25%)	3.9 k (4.23%)	CO	5.0 k (4.52%)	5.0 k (4.31%)	0.9 k (18.00%)
ID	64.4 k (3.60%)	63.5 k (3.64%)	2.8 k (4.42%)	US	4.3 k (3.90%)	5.4 k (4.64%)	3.5 k (65.16%)
BR	52.8 k (2.95%)	47.5 k (2.72%)	9.2 k (19.45%)	TH	3.5 k (3.16%)	3.5 k (2.97%)	0.7 k (20.99%)
BD	45.6 k (2.55%)	46.3 k (2.65%)	3.4 k (7.35%)	AR	3.2 k (2.93%)	3.2 k (2.75%)	0.6 k (18.01%)
IR	40.4 k (2.26%)	40.3 k (2.31%)	2.6 k (6.49%)	PL	2.7 k (2.46%)	2.8 k (2.37%)	0.4 k (15.29%)
TW	26.0 k (1.45%)	25.7 k (1.47%)	1.6k (6.12%)	ZA	2.4 k (2.14%)	2.4 k (2.08%)	0.6 k (23.48%)

We observe a linear correlation between the two metrics of attack intensity (packet rate and bandwidth). However, the number of reflectors does not necessarily correlate with the intensity of attacks. We repeat our analysis for our *primary* dataset and plot it in Fig. 7b. Since this dataset covers a larger number of alerts in 2023, the range of observed attack intensities is also wider compared to the previous case where our data was limited to December 2023. Nevertheless, the same pattern that we discussed earlier, also exists in our *primary* dataset. Note that the number of reflectors in this plot presents a lower bound due to the source address aggregation that we discussed in Section II-A3.

Figure 7b suggests certain DDoS events have an impact outside of the 1500 bytes per packet marker - the typical Maximum Transmission Unit (MTU) for Ethernet. We picked a subset of those alerts and inspected them in detail. We discovered a few explanations for those anomalies. Firstly, it can be attributed to certain attack assets being part of the network. The Extension Mechanisms for DNS (EDNS0) enables DNS server implementations to support large response sizes with 4096 bytes as the recommended default buffer size. Depending on a network's configuration, large DNS responses might not become the subject of IP fragmentation. For another presumable DDoS event, we suspect it to be a data transfer between two networks in which the transfer makes use of jumbo frames. In such a case, configuration thresholds for network anomalies might be surpassed, and a DDoS event was registered.

On the other hand, certain DDoS alerts in Fig. 7b fall below the theoretical lower limit of a DNS packet. We argue that these attacks comprise artificially crafted UDP port 53 traffic and not DNS R&A DDoS traffic.

V. Reflector Demographics

In this section, we analyze the demographics of DNS reflectors that are misused in attacks and compare them to the generic population of open DNS resolvers. For this purpose, we join our *primary* DDoS data with the IP2Location [7] dataset from January to November 2023 to investigate geolocation and network type demographics[2]. We further lookup Autonomous System Numbers (ASNs) of open resolvers using pyasn for the entire 2023.

We limit our analysis to only resolvers that successfully conduct a recursive DNS resolution since these are resolvers with a high potential when misused in attacks. Non-recursive reflectors such as authoritative nameservers that reply with a `Refused` response have instead a minor impact in the DDoS attacks. Also, due to the way the DNS protocol is designed, it is not possible to further limit reflection from these servers. For these reasons, we do not consider them further.

In Table I, we list the top-10 countries hosting open resolvers and compare it to the top-10 countries of open resolvers that are actively misused in DDoS attacks. We observe a surprising bias in the distribution of countries of open resolvers. While China and India together host more than a third (38.8%) of open resolvers in the wild, they do not show up among the top-10 countries of misused open resolvers and only contribute to 3.1% of misused open resolvers together. In the case of China, we also suspect that network monitoring artifacts can play a role in such a concentration of exposed open resolvers. Responses to probing on open resolvers from an external point of view are indistinguishable, and thus we are unable to derive any conclusions on why even the many rather static IPs do not show up on top as misused infrastructure.

This asymmetry raises the question of why, despite having so many open resolvers being hosted in a country, only a small fraction is abused in attacks. Our hypothesis is that the IP churn of open resolvers plays a role in this artifact. In Fig. 8, we explore the correlation between the IP churn of open resolvers per country to their relative misuse. For this purpose, we calculate the average weekly IP churn rate per country by comparing the overlap of IP addresses responding to the weekly scans with the scan results of the previous week over the entire measurement period. For each country, we also calculate the average percentage of misused open resolvers to all open resolvers in that country per week. Confirming that IP churn plays a role, we observe that countries with a low churn on open resolver IPs are more likely to have their infrastructure being abused in attacks, whereas nations with a high churn tend to have a lower presence in the ranking of countries with the most number of misused resolvers. Particularly the IPs

[2] We only have access to the IP2Location dataset during this period.

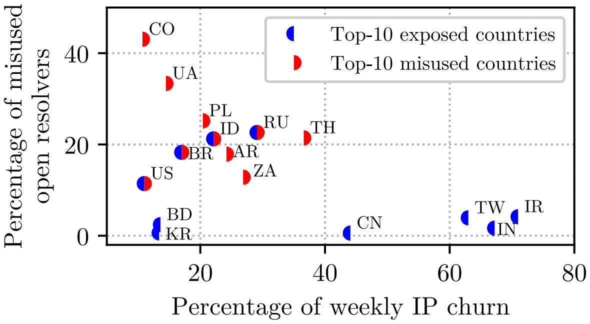

Fig. 8: Comparison of the misuse rate of open resolvers in top-10 countries with the highest number of exposed open resolvers and countries with the highest number of misused open resolvers to the weekly IP churn of open resolvers in the same country. Reflectors in countries with a lower IP churn rate are typically misused more than ones with a high IP churn rate.

of resolvers in China, Taiwan, Iran, and India demonstrate a high churn on the IP assignment while not often being involved in R&A DDoS. In contrast, countries that have a low IP churn tend to rank higher in abused infrastructure. The highest relative misuse is observed in Columbia and Ukraine, comparing top-10 countries with the most exposed and misused open resolvers. These two countries have a relatively low weekly IP churn (∼11% and ∼14%, respectively) for their open resolvers. Having a low IP churn, however, does not guarantee a high misuse. For example, while South Korea and Bangladesh are among countries with a low IP churn, their open resolvers do not frequently appear in DDoS attacks.

We also compare the type of networks in which open resolvers are hosted in the wild to those of resolvers that are misused in the attacks. We observe that the distribution of open resolvers in two categories among various network types are fairly equivalent in values as well as in the ranking of the network types, and thus does not represent an outstanding bias toward a specific network type. This can however be an artifact of aggregation of reflectors among many alerts. Thus, we break down network types per top-20 countries with the most misused open resolvers and compare them with the generic population of open resolvers in these countries. We observe that in most of these countries, misused open resolvers follow a similar distribution to the generic pool of exposed open resolvers, suggesting that misused open resolvers are approximately a random sample of the overall pool. Notable exceptions are India, Iran, and Pakistan, where there is an over-representation of misused open resolvers in data centers.

Additionally, we compared the AS distribution of all open resolvers to misused open resolvers in our primary dataset. All top-10 ASes in both categories are user-access networks, except an educational network among the top-10 networks hosting open resolvers. Comparing the two distributions we observe that only two of the top-10 ASes hosting open resolvers are also in the top-10 list of networks with the highest misuse (ranks #3 and #8 in misused open resolver operators). This reveals a bias toward misusing reflectors in certain networks relative to their generic exposure of open resolvers.

A common feature among the three demographic aspects that we study is that misused open resolvers per aspect have a relatively higher standard deviation compared to the generic exposure of open resolvers. This means that open resolvers in different countries, network types, and ASes are not persistently misused over the year. A potential reason for this could be that different attacker entities rely on different sub-pools of open resolvers to bring about DDoS attacks.

VI. Discussions

In this section, we highlight several recommendations for operators to limit the power of DNS R&A attacks. Although not all of these recommendations stem directly from our research, we believe it is crucial to emphasize them to urge network operators take action. We further discuss the limitations of our study.

A. *Operational Recommendations*

The misuse of open resolvers in DNS-based DDoS attacks has been known for a long time. The problem would not exist if all networks were to implement Source Address Validation (SAV) [4], thus stopping spoofed packets from leaving the network. Nevertheless, the asymmetry between the required effort and benefit of deploying such a filtering has resulted in not all networks being incentivized to do it. This moves the focus for preventing DDoS attacks to the mitigation of open DNS resolvers, which are the second factor facilitating DDoS attacks. However, despite many mitigation efforts and considerable shrinkage in the pool of exposed resolvers, this issue continues to exist and provides attackers with a sizable pool of vulnerable hosts to misuse.

We demonstrate that the untapped potential for DNS reflection is substantially larger than the current number of exploited reflectors. This calls network operators to deploy mechanisms such as ones recommended by the KINDNS [6] initiative to increase the hygiene of their networks and reduce the exposure of open resolvers. Such practices need to be beyond one-off mitigation campaigns due to the dynamicness of the Internet, especially considering that the majority of open resolvers are likely unintentionally exposed as we show in a previous study [32]. Nevertheless, not all open resolvers are unintentional. While we advocate for limiting the access to intentionally deployed recursive DNS servers as much as possible, there might be use cases in which DNS servers need to be publicly accessible (in a similar manner to large public DNS providers). In such scenarios, operators need to consider DNS Response Rate Limiting (RRL) to decrease their contribution to DDoS attacks. Additionally, given the involvement of likely authoritative nameservers in DDoS attacks, it is crucial for their operators to also limit exposure by deploying RRL.

Our results in Section V indicate that a high IP churn can significantly hinder the misuse of open resolvers in DDoS attacks. This might indicate that increasing the IP churn is a solution to reduce the contribution of open resolvers to DDoS attacks. However, we caution network operators against relying on this phenomenon as a primary defense strategy. High IP churn can introduce various operational and network management challenges. Additionally, attackers may adapt by utilizing up-to-date lists of open resolvers, thus negating the impact of high IP churn.

B. Limitations

The core of our work relies on reflectors that are the top contributors to DNS R&A DDoS attacks. This is due to the granularity of the data we have available. As we discuss in Section II-A, we expect this to have a limited impact on our results. A limitation of our work is that we miss insights into the long tail of reflectors. Those however contribute to attacks in a minor manner. We argue therefore that the results we present still provide insight to support, for example, selective take-down or mitigation efforts targeting the most used/most efficient reflectors.

Some considerations should be made about the completeness of the data. The DDoS telemetry dataset used in this paper reports the pool of reflectors that are observed to be participating in attacks. The system has visibility of an attack only if a sensor is on the path of the attack traffic to the victim. This means that our data might account for a lower attack intensity than what the victim experienced in the case in which attack traffic follows a number of paths to the destination. Similarly, an attack could comprise more reflectors than the data are able to report. In addition, we have made the choice of only focusing on single-vector attacks. Despite this, this dataset remains one of the *large-scale* datasets of R&A DDoS attacks available.

Last, while our analysis shed light on possible selection biases in choosing reflectors, it is important to realize that the set of reflectors an attacker might have intended to misuse might be different from the set that is actually misused. For example, some reflectors may become inactive at the time of abuse or they may move to different IP addresses (because of IP churn). This means that indications of patterns or selection biases can only be based, in the scope of this paper, on what we observe in practice.

VII. Related Work

Nawrocki et al. [19] study the DNS amplification ecosystem by leveraging sampled data from a regional IXP and investigate how attackers detect and misuse DNS reflectors. They detect 25.7k attacks over a three-month measurement period. They use packet headers to attribute attack events to an attacking entity. For this entity, they investigate the selection of reflectors showing that attackers continuously update their list of reflectors to misuse. Additionally, they cluster reflectors and report that looking at the similarity of reflectors' sets, attacks are majorly disjoint. They also report that 2% of amplifiers observed at the IXP are authoritative nameservers. Similarly, Kopp et al. [12] analyze traffic flows at a major IXP over seven months and identify roughly 73k DNS amplification DDoS attacks. They report a maximum number of 14k reflectors with 776 reflectors seen on average. Rossow [26] investigates R&A attacks by studying ISP data. Our work is based on DDoS telemetry from distributed nodes. Besides differences in the coverage and scale of studied attacks, we shed light on the diversities of the misused reflectors by comparing them to the entire pool of exposed reflectors.

Multiple papers focus on investigating attacks from honeypots' point of view [5], [13], [22], [29], [30]. While a large enough set of honeypots can ensure detection of the vast majority of ongoing R&A DDoS attacks, they lack insights into the intensity of the attacks as well as the characteristics of the reflectors misused in attacks.

Another line of prior research focuses on characterizing reflectors that are exposed for potential abuse [15], [16], [23], [24], [26], [31], [33], [34]. Our work extends these works by exploring the actual misuse of DNS reflectors in the wild.

Several studies have investigated the ecosystem of booter services and characterized reflectors abused by these services [9]–[11], [14], [27]. Such studies are typically limited to a number of services that can be found on underground fora since there is not a central list of DDoS-for-hire service providers. Additionally, these studies - due to ethical reasons - intentionally limit the scope of their investigation such that they do not majorly contribute to the business of these providers. Our work, on the other hand, focuses on all reflectors visible in DDoS attacks, whether or not they are explicitly abused by booter services. Karami et al. [11] report stable resolvers to be misused more often. However, considering that likely there is a time gap between the discovery of open resolvers and launching DDoS attacks, non-stable reflectors would rather quickly disappear, limiting the observed set to look as being stable. Thus, to have a fair inference one needs to compare open resolvers discovered on a snapshot to attacks that follow with a reasonably short time gap.

VIII. Conclusions

In this paper, we study the DNS R&A DDoS attacks by analyzing $\sim$556 k single-vector DDoS alerts to shed light on the reflector misuse ecosystem.

Our analysis uncovers a notable pattern of misusing open resolvers in specific countries and ASes. However, the same is not true for network types, where we find that the distribution of misused reflectors across different network types closely mirrors the general population of exposed open resolvers.

We also identify a significant number of hosts with open port 53 that are exploited in DDoS attacks, despite not being open recursive resolvers (e.g., authoritative nameservers). While these hosts may not contribute large volumes of traffic in R&A attacks, they pose a potential threat as a resource exhaustion vector. Besides, due to the design of the DNS protocol, they cannot be further mitigated.

Additionally, the selection of abusable infrastructure per country is heavily biased with respect to the observed IP churn. Infrastructure where the country tends to exhibit low IP churn is more likely to be among the list of top misused infrastructure. Similarly, when we observe a high IP churn in a country's open resolvers, it is not part of said list. We conclude that IP churn plays a vital role in the infrastructure that contributes to DDoS attacks.

Furthermore, while at this moment in time attackers seem to rely on a randomly-selected pool of resolvers, our study suggests that the potency of DNS R&A attacks could increase substantially if more targeted selection of reflectors takes place, and by tapping into the pool of underutilized exposed reflectors. Combined with the observation of churning IP addresses, we believe that adversaries with recent scans on abusable infrastructure can leverage even more reflection and amplification potential.

Acknowledgments

This work was partially supported by the GÉANT GN5-1 programme funded by the European Commission.

References

[1] CAIDA, "The UCSD Network Telescope," https://www.caida.org/projects/network_telescope/, 2012.

[2] Censys, https://censys.com, last accessed on 30 August 2024.

[3] D. Dagon, N. Provos, C. P. Lee, and W. Lee, "Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority," in *15th Network and Distributed System Security Symposium (NDSS)*, 2008.

[4] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2827, May 2000.

[5] T. Heinrich, R. R. Obelheiro, and C. A. Maziero, "New Kids on the DRDoS Block: Characterizing Multiprotocol and Carpet Bombing Attacks," in *Passive and Active Measurement Conference*. Springer International Publishing, 2021, pp. 269–283.

[6] ICANN, "KINDNS," https://kindns.org/, 2022.

[7] IP2Location, "IP Address to IP Location and Proxy Information." [Online]. Available: https://www.ip2location.com/

[8] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti, "Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem," in *Proceedings of the Internet Measurement Conference (IMC)*. ACM, 2017, p. 100–113.

[9] M. Karami and D. McCoy, "Rent to Pwn: Analyzing Commodity Booter DDoS Services," *login: the magazine of USENIX & SAGE*, vol. 38, no. 6, pp. 20–23, 2013.

[10] ——, "Understanding the Emerging Threat of DDoS-as-a-Service," in *6th USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET 13)*. USENIX Association, Aug. 2013.

[11] M. Karami, Y. Park, and D. McCoy, "Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services," in *Proceedings of the 25th International Conference on World Wide Web (WWW)*, 2016, p. 1033–1043.

[12] D. Kopp, C. Dietzel, and O. Hohlfeld, "DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks," in *Passive and Active Measurement Conference (PAM)*. Springer-Verlag, 2021, p. 284–301.

[13] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow, "AmpPot: Monitoring and Defending Against Amplification DDoS Attacks," in *Research in Attacks, Intrusions, and Defenses*. Kyoto, Japan: Springer, 2015, pp. 615–636.

[14] J. Krupp, M. Karami, C. Rossow, D. McCoy, and M. Backes, "Linking Amplification DDoS Attacks to Booter Services," in *Research in Attacks, Intrusions, and Defenses*. Springer International Publishing, 2017, pp. 427–449.

[15] M. Kührer, T. Hupperich, J. Bushart, C. Rossow, and T. Holz, "Going Wild: Large-Scale Classification of Open DNS Resolvers," in *Proceedings of the 2015 ACM Internet Measurement Conference - IMC '15*. New York, USA: ACM Press, 2015, pp. 355–368.

[16] M. Kührer, T. Hupperich, C. Rossow, and T. Holz, "Exit from hell? Reducing the Impact of Amplification DDoS Attacks," in *Proceedings of the 23rd USENIX Security Symposium*. San Diego, CA: USENIX Association, 2014, pp. 111–125.

[17] D. Leonard and D. Loguinov, "Demystifying Service Discovery: Implementing an Internet-Wide Scanner," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC)*. ACM, 2010, p. 109–122.

[18] M. Nawrocki, J. Blendin, C. Dietzel, T. C. Schmidt, and M. Wählisch, "Down the Black Hole: Dismantling Operational Practices of BGP Blackholing at IXPs," in *Proceedings of the Internet Measurement Conference*. ACM, 2019, p. 435–448.

[19] M. Nawrocki, M. Jonker, T. C. Schmidt, and M. Wählisch, "The Far Side of DNS Amplification: Tracing the DDoS Attack Ecosystem from the Internet Core," in *Proceedings of the Internet Measurement Conference (IMC)*. ACM, 2021, pp. 419–434.

[20] M. Nawrocki, M. Koch, T. C. Schmidt, and M. Wählisch, "Transparent forwarders: an unnoticed component of the open dns infrastructure," in *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*. ACM, 2021, p. 454–462.

[21] NETSCOUT ASERT Threat Intelligence Team, "DDoS Threat Intelligence Report: Issue 12," https://www.netscout.com/threatreport/wp-content/uploads/2024/04/Threat_Report_2H2023.pdf, 2024, [Online; accessed 2024-04-25].

[22] A. Noroozian, M. Korczyński, C. H. Gañan, D. Makita, K. Yoshioka, and M. van Eeten, "Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service," in *Research in Attacks, Intrusions, and Defenses*. Springer International Publishing, 2016, pp. 368–389.

[23] Y. Nosyk, M. Korczyński, and A. Duda, "Routing Loops as Mega Amplifiers for DNS-Based DDoS Attacks," in *Passive and Active Measurement*. Springer International Publishing, 2022, pp. 629–644.

[24] J. Park, R. Jang, M. Mohaisen, and D. Mohaisen, "A Large-Scale Behavioral Analysis of the Open DNS Resolvers on the Internet," *IEEE/ACM Transactions on Networking*, vol. 30, no. 1, pp. 76–89, 2022.

[25] V. Paxson, "An analysis of using reflectors for distributed denial-of-service attacks," *SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 3, p. 38–47, jul 2001.

[26] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *Proceedings of the 2014 Network and Distributed Systems Security Symposium*. San Diego: Internet Society, 2014, pp. 23–26.

[27] J. J. Santanna, R. van Rijswijk-Deij, R. Hofstede, A. Sperotto, M. Wierbosch, L. Z. Granville, and A. Pras, "Booters — An analysis of DDoS-as-a-service attacks," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 2015, pp. 243–251.

[28] K. Schomp, T. Callahan, M. Rabinovich, and M. Allman, "On Measuring the Client-Side DNS Infrastructure," in *Proceedings of the Internet Measurement Conference (IMC)*. ACM, 2013, p. 77–90.

[29] D. R. Thomas, R. Clayton, and A. R. Beresford, "1000 days of UDP amplification DDoS attacks," in *2017 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 2017, pp. 79–84.

[30] O. van der Toorn, J. Krupp, M. Jonker, R. van Rijswijk - Deij, C. Rossow, and A. Sperotto, "ANYway: Measuring the Amplification DDoS Potential of Domains," in *2021 17th International Conference on Network and Service Management (CNSM)*. IEEE, Oct. 2021, pp. 500–508.

[31] R. Yazdani, A. Hilton, J. van der Ham, R. van Rijswijk-Deij, C. Deccio, A. Sperotto, and M. Jonker, "Mirrors in the Sky: On the Potential of Clouds in DNS Reflection-Based Denial-of-Service Attacks," in *Proceedings of the 25th International Symposium on Research in Attacks, Intrusions and Defenses*, ser. RAID '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 263–275.

[32] R. Yazdani, M. Jonker, and A. Sperotto, "Swamp of Reflectors: Investigating the Ecosystem of Open DNS Resolvers," in *Passive and Active Measurement*, P. Richter, V. Bajpai, and E. Carisimo, Eds. Cham: Springer Nature Switzerland, 2024, pp. 3–18.

[33] R. Yazdani, Y. Nosyk, R. Holz, M. Korczyński, M. Jonker, and A. Sperotto, "Hazardous Echoes: The DNS Resolvers that Should Be Put on Mute," in *7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2023, pp. 1–10.

[34] R. Yazdani, R. van Rijswijk-Deij, M. Jonker, and A. Sperotto, "A Matter of Degree: Characterizing the Amplification Power of Open DNS Resolvers," in *Passive and Active Measurement*. Springer International Publishing, 2022, pp. 293–318.