Swamp of Reflectors: Investigating the Ecosystem of Open DNS Resolvers

Ramin Yazdani(✉), Mattijs Jonker, and Anna Sperotto

Faculty of Electrical Engineering, Mathematics and Computer Science,
University of Twente, Enschede, The Netherlands
{r.yazdani,m.jonker,a.sperotto}@utwente.nl

Abstract. DNS reflection-based DDoS attacks rely on open DNS resolvers to reflect and amplify attack traffic towards victims. While the majority of these resolvers are considered to be open because of misconfiguration, there remains a lot to be learned about the open resolver ecosystem.

In this paper, we investigate and characterize open DNS resolvers from multiple angles. First, we look at indicators that likely suggest an intention behind the existence of open resolvers. To this end, we cross open resolver IP addresses with reverse DNS measurement data and show that a relatively small group of open resolvers unmistakably indicate their service in hostnames (i.e., PTR records). Second, we investigate the extent to which anycast technique is used among open resolvers and show that this is mainly driven by hypergiants. Additionally, we take a look at the exposure of the authoritative nameservers as open recursive resolvers and show that a non-negligible number of authoritative nameservers also serve as open recursors. Finally, we look at the persistency of open resolvers over time. We study open resolvers longitudinally over a three-year period and show that 1% of open resolvers persistently appear in more than 95% of the measurement snapshots.

Keywords: DNS reflection · Open resolvers · DDoS · Amplification

1 Introduction

Our daily lives increasingly rely on digital infrastructure and the Internet. One of the persistent threats against the Internet services are Distributed Denial of Service (DDoS) attacks. In a DDoS attack, an attacker directs a large volume of network traffic toward a victim, which causes disruptions to the legitimate services of the victim by overwhelming its resources or those of its upstream networks. One of the common types of DDoS attacks are Reflection & Amplification (R&A) attacks in which connection-less networking protocols are leveraged to reflect traffic toward DDoS targets. Domain Name System (DNS) is one of the protocols that are typically misused to bring about reflection-based attacks.

DNS reflection is possible by sending spoofed queries (typically) to open DNS resolvers. Leveraging open resolvers as reflectors means that attackers need to

P. Richter et al. (Eds.): PAM 2024, LNCS 14538, pp. 3–18, 2024.
https://doi.org/10.1007/978-3-031-56252-5_1

be in possession of a fresh list of them. This is crucial since open resolvers tend to present a relatively high IP address churn [4,12–14,26,32] and any list of their IP addresses (or at least for a portion of the involved IP addresses) would become outdated rather quickly. For attackers, this means that they need to run discovery scans rather frequently, which might increase the chance that their infrastructure gets detected by network security monitoring systems. As a recent study [8] shows, attackers are selective in scanning networks for exposed services. An alternative solution would involve relying on reflectors that represent a stable behavior, removing the necessity of frequent discovery efforts for attackers. However, it is not yet clear which fraction of the open resolver population is persistent, nor if they are intentionally deployed or what deployment features characterize them.

Open DNS resolvers have been widely studied in the literature [3,4,10,12,13,19–21,26,27,32]. Multiple previous works focus on characterizing open DNS resolvers, either from the aspect of the legitimacy of responses returned by resolvers [4,12,20,21,27] or by investigating the amplification power of these resolvers [19,32]. Our paper complements prior work by characterizing open resolvers from a persistency point of view and investigating indicators of their intentional deployment. Thus, the goal of our research is to investigate differences among open resolvers that indicate the underlying cause of their existence and contribute to their persistence. The main contributions of our paper are that:

- We cross open resolvers discovery data with reverse DNS measurement data and show that (only) a small set of open resolvers likely signal their service through hostnames.
- We perform a longitudinal study on open resolvers and show that the majority of open resolvers are hosted in networks that exhibit persistence in terms of the total number of open resolvers present.
- We investigate the types of networks hosting open resolvers and show that resolvers with indicators of intentional deployment mostly reside in datacenter networks. Besides, we show that the majority of persistent resolvers reside in user access networks.

2 Datasets

2.1 Open Resolver Discovery Scans

We run open resolver discovery scans through which we collected data over a three-year period between October 2020 and October 2023[1] These scans have a frequency of once per week and target all publicly routable IPv4 addresses (in a random permutation) excluding a minority of network prefixes that have requested to be excluded from our scans. For each IP address, we issue a DNS query (toward destination port 53) using a unique subdomain (to avoid caching)

[1] https://research.openresolve.rs.

by embedding the IP address and a timestamp in the query name. If we receive the expected answer to our query, we consider the IP address in question to be an open DNS resolver. These scans provide us with a list of IP addresses that openly do a DNS recursion themselves or forward it to other recursive resolvers. We include transparent forwarders [18] in our list of open DNS resolvers as well. However, we do not capture responses that arrive from unexpected port numbers (i.e., other than port 53) [10]. We plot the number of open resolvers over our three-year measurement period in Fig. 6 (in Appendix A).

2.2 OpenINTEL

We leverage data provided by OpenINTEL [23], which is an active DNS measurement project that measures the DNS state of 63% of the global forward DNS namespace on a daily basis. The measurement is seeded with domain names registered under a large number of gTLDs as well as various ccTLDs. Relevant to our study, OpenINTEL measures the authoritative nameserver records (i.e., `NS`) of domain names and resolves nameserver names to IP addresses (i.e., `A` records), providing us with a picture of the authoritative landscape.

Additionally, OpenINTEL measures the reverse DNS of the entire IPv4 address space daily. In a reverse DNS lookup, an IP address is translated to, among others, a hostname or PTR record (provided a hostname is configured). This can be used for network administration purposes, among others. In this paper, we leverage the reverse DNS measurements to further study the ecosystem of open DNS resolvers.

2.3 IP Anycast

IP anycast [16] is a technique in which the same IP address is shared among multiple devices. Anycast leverages the Border Gateway Protocol (BGP) to route client traffic destined to an anycast IP address to the nearest (preferred) server from the client's perspective. One of the advantages of anycast is that it can reduce latency for remote clients and distribute network traffic over multiple servers.

We leverage Anycast Census data [1,29] to determine if open resolvers reside in anycasted network prefixes. Anycast Census is based on using an anycast testbed network - currently consisting of 20 nodes - to measure anycasted prefixes. It measures one IP address per /24 prefix (which is the minimum globally-routable prefix size) for addresses that are likely to respond to pings. The public dataset of Anycast Census consists of quarterly snapshots.

3 Indicators in PTR Records

Reverse DNS serves a crucial role in network administration and security by mapping IP addresses to hostnames. We begin our investigation by looking at the hostnames of public resolvers to find common patterns that likely suggest an

Table 1. A list of nine popular public DNS resolvers [5], their IPv4 addresses, and hostnames.

Provider	IPv4 address(es)	Hostname(s)	DNS string in hostname	Anycast
Google	8.8.8.8 8.8.4.4	dns.google.	✓	✓
Cloudflare	1.1.1.1 1.0.0.1	one.one.one.one.	✗	✓
Quad9	9.9.9.9 149.112.112.112	dns.quad9.net.	✓	✓
OpenDNS	208.67.222.222 208.67.220.220	dns.opendns.com. dns.umbrella.com. resolver1.opendns.com. resolver2.opendns.com.	✓	✓
NextDNS	45.90.28.0 45.90.30.190	dns1.nextdns.io. dns2.nextdns.io.	✓	✓
CleanBrowsing	185.225.168.168 185.228.169.168	- family-filter-dns2.cleanbrowsing.org.	✓	✓
UltraDNS	64.6.64.6 64.6.65.6 156.154.70.2 156.154.71.2 156.154.70.3 156.154.71.3	rec1pubns1.ultradns.net. rec1pubns2.ultradns.net. - - - -	✓	✓
Yandex	77.88.8.8 77.88.8.1	dns.yandex.ru. secondary.dns.yandex.ru.	✓	✓
OpenNIC	195.10.195.195 51.77.149.139 88.99.98.111 ...	- 139.ip-51-77-149.eu. dns1.dns-ga.de. ...	✓	✗

intention behind their existence. Note that the intention to run a DNS service on a host, however, is not necessarily correlated with running this service openly. We argue that resolvers that are intentionally deployed would likely follow practices similar to public DNS resolvers (i.e., resolvers that are intentionally configured to serve public users) when setting their hostnames. To this end, we consider nine popular public resolvers [5] and the common practices followed by them, which we elaborate on in the following sections. The resolvers under consideration are listed in Table 1.

Out of nine public resolvers, eight resolvers present a similar pattern by embedding the *dns* string in (at least one of) their hostnames. Given the prominent usage of the *dns* string, we consider this a heuristic that likely indicates intentional deployment of an open resolver. However, hostnames do not always reflect services behind individual IP addresses. Many large networks (e.g., cloud providers and user access networks) deploy hostnames in an automated manner by embedding (part of) the IP address of each host as a label under a domain name in their PTR records. To avoid tagging default PTR records that contain *dns* string as an indicator of DNS service running on a specific host, we extend our heuristic by considering two thresholds that we choose experimentally. We tag a /24 prefix with *default PTR setup* if the number of open resolvers that

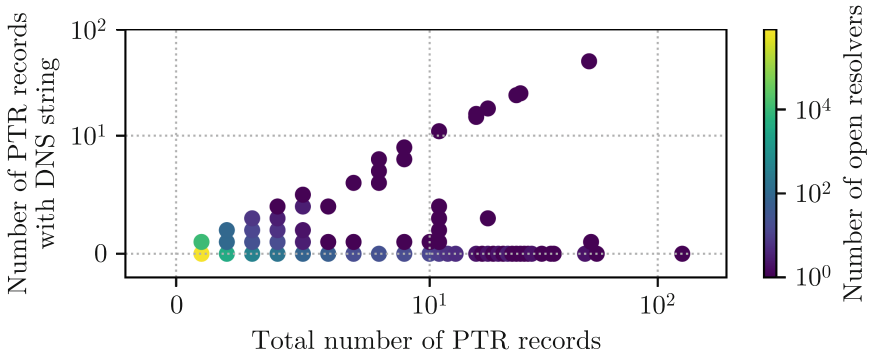

Fig. 1. Number of PTR records (hostnames) including the *dns* string for each open resolver compared to the total number of their PTR records.

have a PTR record containing the *dns* string in that prefix is less than 10% of the hosts in the same /24 prefix with such a string. Furthermore, if the number of PTR records with such a string in that prefix is more than half of the prefix size (127 hosts, indicating a potential default setup), we increase the threshold to 90%. Otherwise, we consider the *dns* string in the hostname of an open resolver to likely be an indicator of its intentional DNS service deployment. We establish these thresholds empirically by investigating their impact on a limited number (i.e., order of tens) of prefixes. Note that our method does not make any inferences on running DNS service openly.

We first investigate the prominence of the *dns* string in the open resolver population by leveraging reverse DNS measurement data from the OpenINTEL project. We observe that the vast majority (∼65.5%) of open resolvers on a single measurement snapshot (2023-Jul-31) do not have a PTR record configured. In Fig. 1 we plot the number of PTR records (hostnames) for each open resolver with at least one hostname compared to the number of (their) PTR records containing the *dns* string. The color of each dot represents the number of resolvers for each data point. We observe two interesting patterns among open resolvers with at least one PTR record. First, a diagonal pattern in which all or the majority of the PTR records contain the *dns* string, with the extreme case being a resolver with 50 PTR records, all containing the *dns* string. We attribute this group to hosts that are more likely to be running DNS service intentionally. The second pattern is a horizontal one in which none of the PTR records contain the *dns* string. We attribute this group to open resolvers that are more likely to be exposed due to misconfigurations. This observation indicates that there is a high variability in both the number of hostnames and the way names are assigned among open resolvers, strengthening our observation that names can assist in identifying intentional deployment.

We then investigate reverse DNS data with respect to network prefixes. On 2023-Jul-31, approximately 1.28 B distinct IP addresses over the entire IPv4 address space have a hostname configured. Roughly 32 M (∼2.5%) of all host-

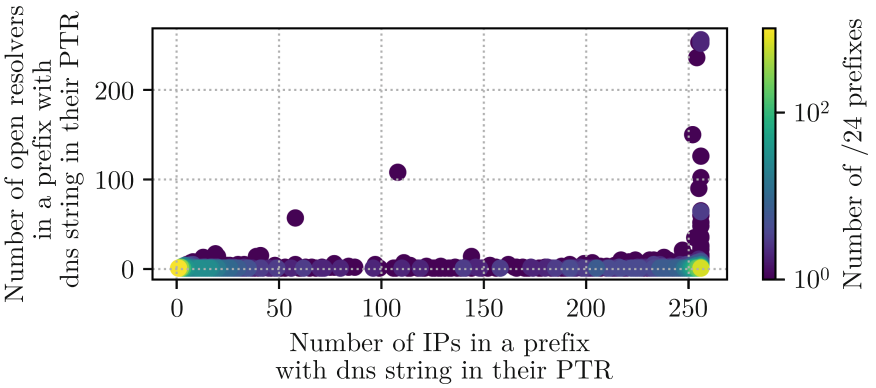

Fig. 2. The scatter plot of the number of IPv4 addresses within a /24 prefix that have *dns* string in their hostnames compared to the number of open resolvers in the same prefix with a similar pattern. The colors show the number of /24 prefixes with the same ratio between the two.

names contain the *dns* string in them. We group all IP addresses with the *dns* string in their hostnames into /24 prefixes. This gives us ∼222 k prefixes. Roughly 51% of these are prefixes in which all of the IP addresses include the *dns* string in their hostnames, likely indicating a default PTR record setup. We then further investigate the reverse DNS data for the open resolvers in the 2023-Jul-31 scan. Out of roughly 2.4 M open resolvers, around 823 k resolvers (∼34.5%) have a hostname configured. Only around 11 k (∼0.46%) resolvers include the *dns* string in at least one of their hostnames, the majority together with their IP address embedded.

For /24 prefixes that host at least one open resolver with *dns* string in its hostname (4.7k prefixes), we plot, per prefix, the number of open resolvers versus the number of IP addresses that embed *dns* string (see Fig. 2). We observe two common patterns in this figure. On the one hand, there is a high number of prefixes in which there are only a handful of PTR records with *dns* string in them as well as a handful of open resolvers with such a string. We argue this group likely indicate their DNS service through their hostnames. On the other hand, there are many networks in which almost the entire block has *dns* string in their hostnames, and in which only a few open resolvers reside. We attribute this group to networks with likely default PTR setups. Alternatively, it is possible that such prefixes are meant to be used as closed resolvers and the limited number of open resolvers in them are misconfigurations. However, we argue that since such networks are typically well-provisioned, misconfigurations in them would be less common compared to CPEs. Using our heuristic with the above-mentioned threshold leaves us with 5.5k (∼0.23%) open resolvers that we suspect to be intentionally running DNS service. Thus, the majority of open resolvers likely do not indicate their service through their PTR records. This suggests that this large portion of resolvers are likely unintentionally open.

Our heuristic will create some false negatives since not all open resolvers necessarily follow such a practice of including the *dns* string in their hostnames. A notable example is Cloudflare public resolver (see Table 1). Additionally, false positives are also possible if a hostname contains the *dns* string as part of an unrelated string[2]. Thus, our results provide an estimate of the number of open resolvers that are intentionally running DNS service.

Key Takeaway: *Although PTR record registration is not mandatory, looking at the open resolvers that have a hostname configured, we observe that the vast majority do not follow a similar pattern to popular public DNS providers. This suggests that most open resolvers are likely unintentionally deployed.*

4 Anycasted Open Resolvers

A common feature among popular public resolvers is to use IP anycast to improve performance for clients in various regions. As seen in Table 1, all but one of the public resolvers in our study are anycasted. The only provider that does not use anycast is OpenNIC, which uses volunteer-provided DNS resolvers. Sommese et al. [28] studied anycast adoption in the authoritative nameserver ecosystem and proved that it is mainly driven by hypergiants, as this is a complex and costly technology. Intuitively, we expect the hypergiants to be the main contributors to the anycasted recursive resolver ecosystem.

We investigate the extent to which open resolvers are located in anycasted networks. We fuse the open resolver scan data (2022-Jan-17) with anycast census data [1,29] from the same month and observe that roughly 9.3 k (∼0.34%) resolvers are hosted in 216 distinct `/24` anycasted network prefixes. In terms of Autonomous Systems (ASes), this translates to 58 distinct networks. In Table 2 we provide the top ASes hosting anycasted open resolvers. Cloudflare and Akamai together are behind 56% of anycasted open resolvers. Additionally, as one expects, multiple public DNS providers are visible in this table. Looking at the network types, 99% of the anycasted resolvers are hosted in DCH (datacenter/hosting/transit) and CDN networks which are expected network types for anycasted infrastructure. Those findings confirm our intuition that hypergiants are prominently present in the open resolver ecosystem, but the same findings also show that anycast is only minimally present among open resolvers.

Key Takeaway: *While popular public DNS providers hold a large share of the anycasted recursive DNS infrastructure, using anycast is marginal among the overall crowd of open resolvers.*

5 Authoritative Open Resolvers

Authoritative nameservers are meant to serve queries for specific subdomains, i.e., zones. However, there is a possibility to combine an authoritative nameserver

[2] A manual inspection shows that this is not common in our dataset.

Table 2. The top-10 origin ASes for open resolvers in anycasted prefixes.

ASN	AS Name	# Resolvers	% Resolvers
AS13335	Cloudflare, Inc.	2636	28.5%
AS21342	Akamai International B.V.	2550	27.5%
AS23393	NuCDN LLC	1022	11.0%
AS205157	Daniel Cid	1021	11.0%
AS34939	NextDNS, Inc.	1020	11.0%
AS45102	Alibaba (US) Technology Co., Ltd.	510	5.5%
AS397213	Neustar Security Services	108	1.2%
AS397232	Neustar Security Services	107	1.2%
AS398962	CONTROLD INC.	32	0.3%
AS3549	Level 3 Communications, Inc.	23	0.2%

with a recursive nameserver, despite the fact that this is not generally considered a good practice [6]. For example, *BIND* DNS software [2] allows configuring a server to act simultaneously as an authoritative and a recursor. This means that there is a possibility for authoritative nameservers to be exposed as open resolvers.

To investigate the degree to which authoritative nameservers contribute to the open resolvers ecosystem, we leverage OpenINTEL authoritative nameserver measurement data and fuse it with the open resolvers discovery scans. We study all domains in *.com*, *.net*, *.org* Top-Level Domains (TLDs) as well as all new Generic Top-Level Domains (gTLDs) available through ICANN's CZDS. We create a mapping of second-level domain names in these zones to the IP address of their nameservers. However, domain names may point their nameservers to open resolvers while the resolver is not necessarily the authoritative server for that domain. To investigate the extent to which open resolvers coincide with authoritative nameservers, we issue DNS queries for domain names in our list and send these queries to their open resolver nameserver addresses. If an open resolver is in fact authoritative for a domain name, the DNS answer should have the `AA` (authoritative answer) flag set. However, it is known [21] that open resolvers might even set the AA flag when they are not authoritative for the domain name in question. Looking at the open resolver discovery scans on 2023-Oct-30 we find 10.6 k of these non-compliant open resolvers. We exclude these resolvers from our study since it is not possible to infer if they are also running an authoritative nameserver using the DNS `AA` flag.

The mapping of second-level domains to their nameservers gives us 666 M distinct *{domain name, IP address}* tuples. Next, we cross this mapping with the (`AA` flag compliant) open resolvers list to get a list of domain names for which the nameservers happen to be authoritative and an open resolver. This results in roughly 468 k *{domain name, IP address}* tuples. This involves 248 k unique

Table 3. The top-5 TLDs for domain names that have at least one nameserver address that is authoritative and an open resolver.

TLD	# Domains	% Domains
com	205.9 k	83.1%
net	19.6 k	7.9%
org	10.7 k	4.3%
vip	0.9 k	0.4%
app	0.9 k	0.3%

domain names pointing their NS addresses to 11.8 k unique open resolvers. In Table 3 we summarize the details of our dataset.

Our queries result in 441 k answers concerning 246 k unique domain names. 238 k of these domains have the `AA` flag set in at least one of the answers from their authoritative nameservers. Approximately 18 k domains have at least one answer coming from a server which is not authoritative for that domain name. Around 10 k domain names use a mix of authoritative and non-authoritative nameservers. Our results indicate that the majority (∼95%) of open resolvers that are listed as an authoritative nameserver for a domain name are indeed running the recursive and authoritative services simultaneously. This could constitute a potential security risk since authoritative nameservers are typically well-provisioned and this could make them attractive for being misused in R&A DDoS attacks.

Our methodology to infer the open resolvers that are serving as an authoritative nameserver provides a lower boundary to this number, for two reasons. First, our dataset includes only second-level domains. Second, we cover a subset of TLDs in our measurement. However, it avoids false positives that occur by purely relying on PTR records to infer infrastructure resolvers [22].

Key Takeaway: *Authoritative nameservers are typically deployed with rich resources. Co-hosting these servers with an open resolver might constitute a higher risk of misuse in R&A DDoS attacks.*

6 Persistency over Time

We postulate that resolvers that are intentionally deployed as open should present a persistent availability over time compared to the lifetime of resolvers that are unintentionally exposed. In order to investigate the persistency of open resolvers over time, we look at the longitudinal open resolvers dataset (154 snapshots in three years) and calculate the percentage of snapshots in which each resolver appears among the snapshots that follow its initial discovery. We only investigate resolvers in the first year of snapshots to make sure that the persistency of resolvers is fairly examined.

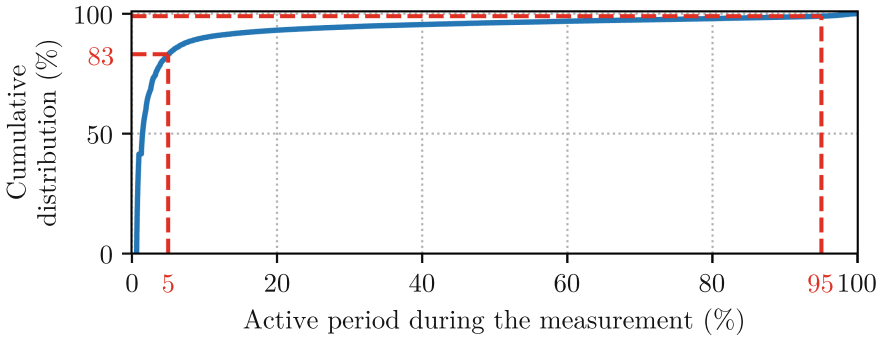

Fig. 3. The distribution of the relative presence of open resolvers during the measurement period.

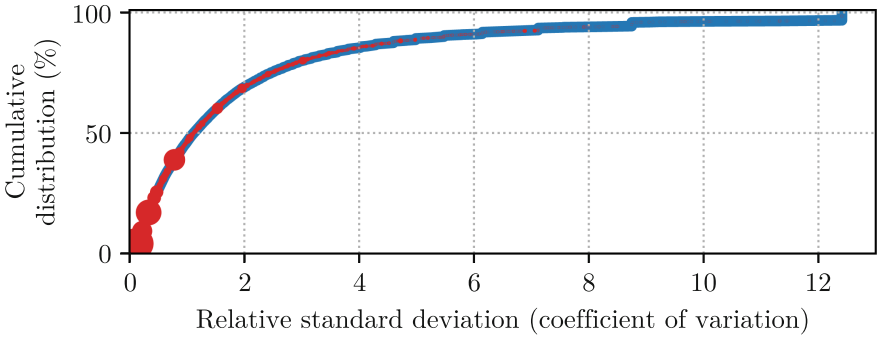

Fig. 4. The distribution of the relative standard deviation (σ/μ) for the number of open resolvers per AS (blue line). We also map open resolvers on a single snapshot to their ASes (red circles). (Color figure online)

We observe roughly 32 M distinct open resolver IPv4 addresses over the first year of our measurement. This number is much higher than the average number of open resolvers (roughly 2.7 M resolvers) per snapshot due to the dynamic appearance of open resolvers (e.g., due to IP churn). Roughly 83% of these IP addresses are visible on less than 5% of our snapshots, while 1% are consistently present in more than 95% of the snapshots (see Fig. 3).

We also look at the persistency of ASes that host open resolvers. To this end, we calculate the Relative Standard Deviation (RSD) of the number of open resolvers per AS and plot a CDF in Fig. 4 (blue line). The RSD (also known as the coefficient of variation) is defined as the ratio of the standard deviation (σ) to the mean value (μ). A lower relative standard deviation means that the number of open resolvers in a network does not change much over time. We also map open resolvers we observed during a single measurement snapshot to their ASes. We plot these as red circles in Fig. 4, where the circle diameter is proportional to the number of resolvers. Figure 4 shows that the majority of open resolvers are hosted in ASes that have a lower RSD, meaning that those ASes show a

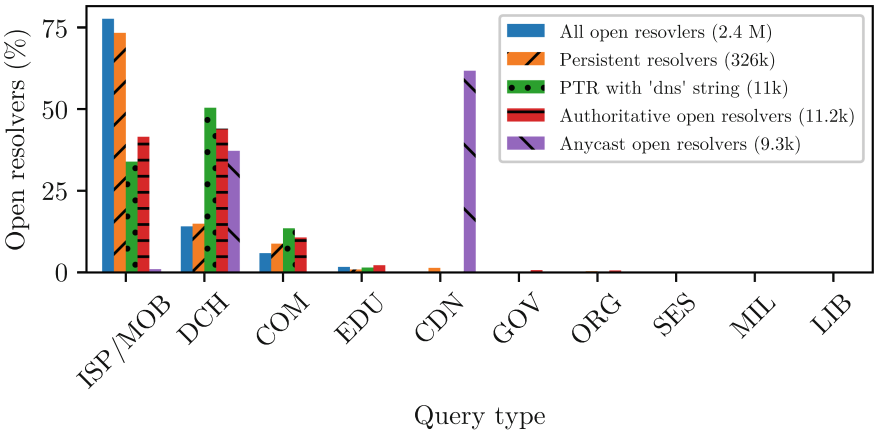

Fig. 5. The distribution of open resolvers in different network categories.

persistent behavior over time in terms of the number of open resolvers in their network.

The persistent concentration of open resolvers in a limited number of networks could have repercussions on the misuse of these servers in R&A attacks, as it could imply that an attacker might not need to sweep the entire address space to find reflectors.

Key Takeaway: *The majority of open resolvers persistently reside in ASes that have a stable number of open resolvers. We argue that this could constitute a security risk, but also it holds potential for prioritized mitigation efforts.*

7 Network Types

Investigating the type of networks hosting open DNS resolvers has been conducted in earlier work. Kührer et al. [12] analyze the reverse DNS records of short-lived open resolvers and attribute them to residential networks. Schomp et al. [26] take a different approach and rely on various metrics such as HTTP probes to infer open resolvers in residential networks. We leverage the IP intelligence data of IP2Location [7] to attribute open resolvers to various network types and corroborate the findings of earlier work. The added value of our method is that it allows us to categorize the network types of all open resolvers even if e.g., they do not set a PTR record.

Looking at the types of the networks that host open resolvers (see Fig. 5), we observe that eyeball networks host the majority (∼80%) of general open resolvers, while open resolvers with *dns* string in their hostname(s) are mainly (∼50%) hosted in datacenters. Considering the resolvers that are persistently available on more than 95% of the scan snapshots, we see a similar distribution to the generic population of open resolvers, with 72% of these resolvers residing in user access networks. The large share of persistent resolvers in ISP networks

is potentially because of providers with a static IP address assignment, even though these resolvers are likely not intentionally deployed.

From a DDoS attacker point of view, there is a trade-off between misusing different categories of open resolvers. While datacenter-based resolvers benefit from a high link-capacity [30], these resolvers are often considered to deploy measures that limit damage in case they are misused in DDoS attacks. Additionally, there is a chance that these resolvers are honeypots deployed to detect malicious network activities. On the other hand, leveraging ISP-based resolvers has a lower risk for attackers, while these resolvers typically have a lower link-capacity compared to datacenter-based resolvers.

Key Takeaway: *Persistent open resolvers are mostly hosted in user access networks. This, combined with the intuition that such resolvers likely exist due to misconfigurations, raises worries about their potential misuse in attacks.*

8 Related Work

Open DNS resolvers have been considered a threat to the Internet for a long time by majorly being involved in DDoS attacks [9,13,17,25] and DNS cache poisoning attacks [11]. Although the number of open resolvers has reduced substantially over time [12,21,32], the remaining ones are multiple orders of magnitude more than what attackers typically misuse in DDoS attacks [17].

Kührer et al. [12] classify the DNS software of open resolvers as well as the hardware of their underlying. Besides, they study the IP churn of open resolvers over time and (by investigating the reverse DNS PTR records) report that a big part of the dynamicity of open resolvers is due to dynamic IP connections. Finally, they look at the authenticity of the answers provided by open resolvers and investigate the intentions behind DNS response manipulation. Similarly, Park et al. [20,21] investigate the authenticity of DNS responses returned by open resolvers and the intentions behind these manipulated responses. Considering the focus of our work being DDoS attacks, we only investigate open resolvers that return an expected answer and thus are more appealing for DDoS attackers.

Doan et al. [5] investigate the performance of popular public DNS resolvers using the RIPE Atlas platform [24]. Although our work does not focus on DNS performance, we base our analysis on the common practices followed by the public resolvers in this work.

Several studies focus on the DDoS potential that open DNS resolvers expose. Leverett and Kaplan [15] estimate a lower bound for the global R&A DDoS potential by leveraging the network speed measurements from the M-Lab project. Nosyk et al. [19] study the packet amplification potential of DNS queries and attribute them to routing loops and middleboxes. Similarly, Yazdani et al. [31] investigate the underlying causes of DNS packet amplification and report directed IP broadcasting to be an additional underlying cause for this phenomenon. The bandwidth amplification power of open resolvers was studied by Yazdani et al. [32], showing that there is a difference in the extent to which open resolvers can return amplified DNS responses.

A common practice among the existing work is that they do not differentiate deliberately public DNS resolvers from open resolvers that are exposed due to misconfigurations. To the best of our knowledge, our work is the first study that investigates the indicators of intentional deployment of open resolvers and sheds light on their dynamics.

9 Ethical Considerations

Our study involves active DNS measurements to discover open resolvers. This introduces some ethical concerns that we have carefully considered when designing our measurement setup. First, we limit the probes that we send to each distinct IP address to one per week. Additionally, the probes were randomized in the IPv4 address space so that remote networks would not receive bursts of queries. We also indicate the purpose of our queries - by using a domain name under our control - both in the PTR record of our scanner machine as well as in the queries that we issue. Finally, we run a Web server for this domain name on which we explain how network operators can opt out of our study.

Additionally, we run measurements to investigate open resolvers that are authoritative for one or more domain names. We follow common good practices in conducting these measurements. For example, we randomize queries for distinct domain names as well as for distinct authoritative nameservers, so that we do not disrupt the tested networks.

10 Conclusions

DDoS attackers typically rely on open resolvers to direct a large amount of network traffic towards victims. Open resolvers, however, present some diversity due to differences in the underlying cause of their existence. Although little is known about the strategies that attackers take to select their set of resolvers to abuse, such diversities can be leveraged to optimize the amount of traffic that is sent to a victim. Our results show that only 0.23% of open resolvers likely indicate an intentional deployment, meaning that there is still ample opportunity to reduce the number of DNS reflectors. Our research also shows that specific groups of open resolvers, namely open resolvers co-hosted with authoritative nameservers and persistent resolvers in user access networks, could enable prioritized mitigation efforts.

Acknowledgments. We would like to thank the anonymous reviewers and our shepherd Ramakrishna Padmanabhan for their valuable feedback on our paper. We also thank Raffaele Sommese for his insightful feedback in the early stages of this research. We gratefully acknowledge the support by ip2location.com, who provided us with an academic license to use their data. This work was partially supported by the GÉANT GN5-1 programme funded by the European Commission.

A Number of Open Resolvers over Time

In Fig. 6 we plot the number of open resolvers that we observe in our scans. The dip in January 2022 is likely due to a measurement error.

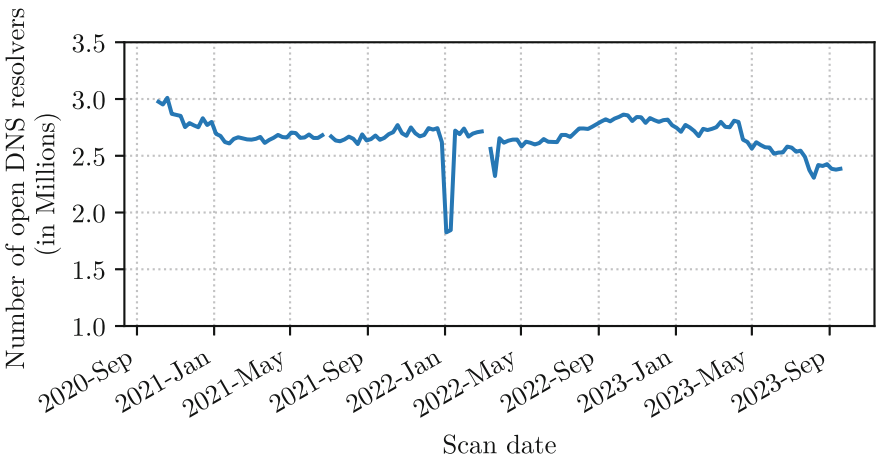

Fig. 6. Number of open DNS resolvers observed in our weekly scans between October 2020 and October 2023.

References

1. Anycast Census, a joint project of the University of Twente, SIDN, and CAIDA. https://github.com/ut-dacs/Anycast-Census/
2. BIND Dns. https://www.isc.org/bind/
3. Al-Dalky, R., Rabinovich, M., Schomp, K.: A look at the ecs behavior of dns resolvers. In: Proceedings of the Internet Measurement Conference, pp. 116–129. IMC '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3355369.3355586
4. Dagon, D., Provos, N., Lee, C.P., Lee, W.: Corrupted DNS resolution paths: the rise of a malicious resolution authority. In: 15th Network and Distributed System Security Symposium (NDSS) (2008)

5. Doan, T.V., Fries, J., Bajpai, V.: Evaluating public DNS services in the wake of increasing centralization of DNS. In: 2021 IFIP Networking Conference (IFIP Networking), pp. 1–9 (2021). https://doi.org/10.23919/IFIPNetworking52078.2021.9472831
6. ICANN: KINDNS (2022). https://kindns.org/guidelines/public-resolver-operators/
7. IP2Location: IP Address to IP Location and Proxy Information. https://www.ip2location.com/
8. Izhikevich, L., Tran, M., Kallitsis, M., Fass, A., Durumeric, Z.: Cloud watching: understanding attacks against cloud-hosted services. In: Proceedings of the 2023 ACM on Internet Measurement Conference, pp. 313–327. IMC '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3618257.3624818
9. Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., Dainotti, A.: Millions of targets under atack: a macroscopic characterization of the DoS ecosystem. In: Proceedings of the ACM SIGCOMM Internet Measurement Conference. vol. Part F131937, pp. 100–113 (2017). https://doi.org/10.1145/3131365.3131383
10. Kaizer, A.J., Gupta, M.: Open resolvers: understanding the origins of anomalous open DNS resolvers. In: Mirkovic, J., Liu, Y. (eds.) Passive and Active Measurement, pp. 3–14. Springer International Publishing, Cham (2015)
11. Kaminsky, D.: Black Ops 2008: It's The End Of The Cache As We Know It. Black Hat USA (2008)
12. Kührer, M., Hupperich, T., Bushart, J., Rossow, C., Holz, T.: Going wild: large-scale classification of open DNS resolvers. In: Proceedings of the 2015 ACM Internet Measurement Conference - IMC '15, pp. 355–368. ACM Press, New York, USA (2015). https://doi.org/10.1145/2815675.2815683
13. Kührer, M., Hupperich, T., Rossow, C., Holz, T.: Exit from hell? reducing the impact of amplification DDoS attacks. In: Proceedings of the 23rd USENIX Security Symposium, pp. 111–125 (2014)
14. Leonard, D., Loguinov, D.: Demystifying service discovery: implementing an internet-wide scanner. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, pp. 109–122. IMC '10, Association for Computing Machinery, New York, NY, USA (2010). https://doi.org/10.1145/1879141.1879156
15. Leverett, E., Kaplan, A.: Towards estimating the untapped potential: a global malicious DDoS mean capacity estimate. J. Cyber Policy **2**(2), 195–208 (2017). https://doi.org/10.1080/23738871.2017.1362020
16. Mendez, T., Milliken, W., Partridge, D.C.: Host Anycasting Service. RFC 1546 (Nov 1993). https://doi.org/10.17487/RFC1546, https://www.rfc-editor.org/info/rfc1546
17. Nawrocki, M., Jonker, M., Schmidt, T.C., Wählisch, M.: The far side of DNS amplification: tracing the DDoS attack ecosystem from the internet core. In: Proceedings of the 21st ACM Internet Measurement Conference, pp. 419–434. IMC '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3487552.3487835
18. Nawrocki, M., Koch, M., Schmidt, T.C., Wählisch, M.: Transparent forwarders: an unnoticed component of the open DNS infrastructure. In: Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies, pp. 454–462. CoNEXT '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3485983.3494872

19. Nosyk, Y., Korczyński, M., Duda, A.: Routing loops as mega amplifiers for DNS-based DDoS attacks. In: Hohlfeld, O., Moura, G., Pelsser, C. (eds.) Passive and Active Measurement, pp. 629–644. Springer International Publishing, Cham (2022)
20. Park, J., Jang, R., Mohaisen, M., Mohaisen, D.: A large-scale behavioral analysis of the open DNS resolvers on the internet. IEEE/ACM Trans. Network. **30**(1), 76–89 (2022). https://doi.org/10.1109/TNET.2021.3105599
21. Park, J., Khormali, A., Mohaisen, M., Mohaisen, A.: Where are you taking me? behavioral analysis of open DNS resolvers. In: 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2019, pp. 493–504. IEEE (2019). https://doi.org/10.1109/DSN.2019.00057
22. Pearce, P., et al.: Global measurement of DNS manipulation. In: 26th USENIX Security Symposium (USENIX Security 17), pp. 307–323. USENIX Association, Vancouver, BC (Aug 2017), https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/pearce
23. van Rijswijk-Deij, R., Jonker, M., Sperotto, A., Pras, A.: A high-performance, scalable infrastructure for large-scale active DNS measurements. IEEE J. Sel. Areas Commun. **34**(6), 1877–1888 (2016). https://doi.org/10.1109/JSAC.2016.2558918
24. RIPE NCC: RIPE Atlas: A global internet measurement network. Internet Protocol J. **18**(3), 2–26 (2015). http://ipj.dreamhosters.com/wp-content/uploads/2015/10/ipj18.3.pdf
25. Rossow, C.: Amplification hell: revisiting network protocols for DDoS abuse. In: Proceedings of the 2014 Network and Distributed Systems Security Symposium, pp. 23–26. No. February, Internet Society, San Diego (2014). https://doi.org/10.14722/ndss.2014.23233, http://www.internetsociety.org/sites/default/files/01_5.pdf
26. Schomp, K., Callahan, T., Rabinovich, M., Allman, M.: On measuring the client-side DNS infrastructure. In: Proceedings of the 2013 Conference on Internet Measurement Conference, pp. 77–90. IMC '13, Association for Computing Machinery, New York, NY, USA (2013). https://doi.org/10.1145/2504730.2504734
27. Schomp, K., Callahan, T., Rabinovich, M., Allman, M.: Assessing DNS vulnerability to record injection. In: Faloutsos, M., Kuzmanovic, A. (eds.) Passive and Active Measurement, pp. 214–223. Springer International Publishing, Cham (2014)
28. Sommese, R., et al.: Characterization of anycast adoption in the DNS authoritative infrastructure. In: Network Traffic Measurement and Analysis Conference (TMA'21) (2021)
29. Sommese, R., et al.: MAnycast2: using anycast to measure anycast. In: Proceedings of the ACM Internet Measurement Conference, pp. 456–463. IMC '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3419394.3423646
30. Yazdani, R., et al.: 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, pp. 263–275. RAID '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3545948.3545959
31. Yazdani, R., Nosyk, Y., Holz, R., Korczyński, M., Jonker, M., Sperotto, A.: Hazardous echoes: the dns resolvers that should be put on mute. In: 2023 7th Network Traffic Measurement and Analysis Conference (TMA), pp. 1–10 (2023). https://doi.org/10.23919/tma58422.2023.10198955
32. Yazdani, R., van Rijswijk-Deij, R., Jonker, M., Sperotto, A.: A matter of degree: characterizing the amplification power of open DNS resolvers. In: Hohlfeld, O., Moura, G., Pelsser, C. (eds.) PAM 2022. LNCS, vol. 13210, pp. 293–318. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-98785-5_13