

# Routing Loops as Mega Amplifiers for DNS-based DDoS Attacks

Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France  
{first.last}@univ-grenoble-alpes.fr

**Abstract.** DDoS attacks are one of the biggest threats to the modern Internet as their magnitude is constantly increasing. They are highly effective because of the amplification and reflection potential of different Internet protocols. In this paper, we show how a single DNS query triggers a response packet flood to the query source, possibly because of middleboxes located in networks with routing loops. We send DNS A requests to 3 billion routable IPv4 hosts and find 15,909 query destinations from 1,742 autonomous systems that trigger up to 46.7 million repeating responses. We perform traceroute measurements towards destination hosts that resulted in the highest amplification, locate 115 routing loops on the way, and notify corresponding network operators. Finally, we analyze two years of historical scan data and find that such “mega amplifiers” are prevalent. In the worst case, a single DNS A request triggered 655 million responses, all returned to a single host.

**Keywords:** DDoS · DNS resolvers · Amplification attacks · Reflection attacks · Routing loops

## 1 Introduction

Distributed Denial-of-Service (DDoS) attacks have become increasingly common and constantly growing in size. One of the largest known attacks against Google services already peaked at 2.54 Tbps and the attack volume is likely to get more important with time [18]. The two main factors that contribute to the effectiveness of DDoS attacks are *reflection* and *amplification*. Attackers use Internet services that satisfy two requirements: respond to their requests (*reflect*) and generate either a large number of responses or a response of a much larger size (*amplify*) towards a victim. Reflection attacks are only effective when compromised hosts (bots) send requests with spoofed IP addresses. Consequently, they need to be located in networks that do not deploy Source Address Validation (SAV), known as Best Current Practice 38 (BCP-38) [22, 50], for outgoing traffic.

Several initiatives aim at reducing the possibility of DDoS attacks [1, 5, 6, 11, 27, 28, 30, 34, 37, 38, 46, 47, 49, 55, 61], for instance, measurements of the amplification potential of different protocols and notifications of the affected parties. Other non-profit initiatives, such as Shadowserver Foundation [51], provide daily

reports to network operators and 132 national Computer Security Incident Response Teams (CSIRTs).

Amplifying services are mostly UDP-based because of their connectionless nature. An attacker sends spoofed requests and the services reflect responses to victims. The most prominent UDP reflectors are NTP and DNS [20, 21], which have been leveraged by several attack vectors [1, 6, 11, 34, 46]. Theoretically, the TCP three-way handshake prevents the connection establishment with spoofed hosts because the response of the reflecting service goes to the victim and not to the host launching the attack. Nevertheless, certain TCP implementations are prone to amplification [27, 28] and potentially with infinite amplification factors [5].

In the concurrent work, Bock *et al.* [5] located middleboxes inside routing loops by sending a sequence of carefully crafted TCP packets. They even found 19 IP addresses that triggered infinite loops. In our work, we show that a trivial DNS A request is enough to trigger a similar behavior. Moreover, we identify 64 IP addresses triggering possibly infinite amplification. Our methodology consists of probing the whole routable IPv4 address space with DNS A requests to find 15,909 destination addresses from 1,742 autonomous systems (ASes) triggering up to 46.7 million identical response packets. We then run traceroute measurements towards 435 destination hosts that resulted in the highest amplification and identify 115 routing loops involving 35 autonomous systems. We have reported these findings to network operators. Finally, we analyze 2 years of packet traces from our DNS scans in both IPv4 and IPv6 address spaces to find 944,087 requests that triggered repeating responses—397 of them caused more than 1,000 responses and 18 requests caused more than 1 million responses. As an extreme case, one DNS A request triggered 655 *million* responses.

The rest of the paper is organized as follows. Section 2 provides background on DDoS attacks, amplification, and reflection. Section 3 describes the threat model and Section 4 introduces the measurement setup. We present scan results and analyze the persistence of the vulnerability in Section 5. Section 6 discusses ethical considerations and disclosure. Finally, we present related work in Section 7 and conclude in Section 8.

## 2 Background on DDoS Attacks

One of the largest DDoS attacks known to date took place in September 2017 and was reported by Google in October 2020 [18]. Attackers sent spoofed requests to SNMP, CLDAP, and DNS servers that, in turn, sent amplified responses to Google. The reflected traffic peaked at 2.54 Tbps. In February 2020, Amazon Web Services (AWS) reported an attack using hijacked CLDAP servers that generated traffic up to 2.3 Tbps [3]. If measured in requests per second (rps), two prominent attacks happened in 2021: Yandex [45] and Cloudflare [62] reported receiving 21.8 million and 17.2 million rps, respectively. As the Internet grows in terms of computing power, bandwidth, and the number of connected devices, the volume of DDoS attacks becomes increasingly high [18].

A DDoS attack aims to overwhelm the victim service with a tremendous amount of traffic to prevent legitimate clients from using the service. Although an attacker alone may achieve this effect, large-scale attacks usually rely on botnets, networks of compromised machines that receive instructions from the command-and-control (C&C) center (operated by the attacker).

The real danger of DDoS attacks comes from *reflectors* and *amplifiers*. A *reflector* is a machine that accepts a request (with a spoofed source IP address) and sends a response [44]. There are millions of services on the Internet such as web servers or open DNS resolvers that can act as reflectors. Once reflecting services are located, the attacker instructs the botnet under her/his control to start sending requests. Requests with the spoofed source IP addresses of victims are sent to reflectors. As a result, the victim receives all the reflected traffic. Carefully crafted requests can trigger reflectors to send large or numerous responses; such reflectors are called *amplifiers*.


There are several ways to assess the effectiveness of a DDoS attack. We can measure the absolute amount of generated traffic in packets per second (pps), bits per second (bps) or requests per second (rps) [18]. In the case of amplification, another informative metric is the ratio of traffic generated by the amplifier to the traffic needed to trigger the amplifier. Rossow [47] proposed two units of measurement: bandwidth amplification factor (BAF) and packet amplification factor (PAF). BAF divides the size of the packet payload sent from the amplifier to the victim by the size of the packet payload sent from the attacker to the amplifier. Likewise, PAF divides the number of packets sent to the victim by the number of packets sent to the amplifier. In both cases, the higher the value, the more destructive the attack is.

In the remainder of this paper, we use Rossow's [47] packet amplification factor (PAF) metric to assess the amplification potential of DNS queries caught in routing loops. As we only send one DNS request, the PAF is always equal to the number of received responses.

### 3 Threat Model

Our threat model is an amplified and reflective DDoS attack in which the attacker sends DNS queries. Therefore, we first recall how regular DNS resolution operates. It starts with a client sending its DNS request to a recursive resolver. This entity is capable of following the domain name tree from the root down to the authoritative nameservers of a given domain. Recursive resolvers heavily rely on caching and query prefetching to speed up the resolution process. Whether it succeeds or not, a recursive resolver returns a response packet to the client with one of the defined response codes [35]. Thus, a client expects to receive a single response packet for a single request.

When one is constantly receiving multiple copies of the same packet, there might be some routing anomaly on the way between the sender and the receiver, such as loops. Routing loops are a well-known, old phenomenon, extensively studied in the literature [19, 32, 43, 54, 60, 63–65]. They fall into two broad cate-




**Fig. 1.** The sender (1.2.3.4) initiates a request to the receiver (5.6.7.8). The packet travels through R1 and R2 until being caught in the loop involving R2, R3, R4. Although the request never reaches the receiver, the sender receives replies from the middlebox.

gories: transient and persistent. Transient loops appear when topology changes and the routing protocol has not yet converged. Such loops do not require manual intervention to be resolved. Persistent loops are likely to be a result of a misconfiguration, such as announcing addresses that are routable but not allocated [32, 60]. A packet entering the routing loop is very likely not to reach the destination. Xia *et al.* [60] analyzed the location of routing loops and found that the majority of them involve destination autonomous systems. We report similar findings later in Section 5.2. Consequently, the same loops can be triggered from multiple vantage points.

Recently, Bock *et al.* [5] discovered that networking middleboxes (such as firewalls or national censors), when located inside the routing loop, continuously process a request caught in a loop and keep responding to it. Figure 1 illustrates such a setup. The sender (1.2.3.4) sends a request to the receiver (5.6.7.8) via R1 and R2. Somewhere on its way (in transit or at the destination autonomous system), the request packet enters the routing loop between R2, R3, and R4. Each time the packet goes from R2 to R3, it triggers the middlebox to respond to the sender (more precisely, to the host with the source IP address of the packet, which can be spoofed). The receiver (5.6.7.8) never sees the request. Such a looping packet should be dropped when its time-to-live (TTL) reaches 0. However, if the TTL is not decreased for any reason, the packet may loop infinitely (or until a reboot or router failure drops it).

An attacker knowing about the presence of routing loops and middleboxes can achieve two principal goals: saturate links involved in the routing loop and reflect the generated responses. If the loop is located in the destination autonomous system (AS) and the spoofed source IP address belongs to the same AS, such a packet may be dropped at the network edge even before reaching the loop. It happens when SAV for incoming traffic drops the packet from the outside with the source IP belonging to the inner network. However, recent work showed that inbound SAV is not widely deployed [25, 12, 23, 24].

We have very few assumptions about the capabilities of the attacker. Most importantly, (s)he has to be located in the network that allows outbound spoof-



**Fig. 2.** Measurement setup for the DNS scan. The scanner sends a DNS A request to the recursive resolver (1.2.3.4). The resolver contacts the authoritative nameserver, obtains the response, and returns it to the scanner.

ing. Thus, DNS packets with spoofed IP addresses can leave the network. Recent work showed that such misconfigured networks are still not uncommon on the Internet [33, 31] and they are publicly listed [8]. The attacker does not have any special hardware or software requirements, because a single DNS packet, occasionally resent, is enough to keep the loop going. Finally, it is not necessary to register a domain name as any existing one can be queried.

## 4 Internet-Wide Scans

Our measurement technique relies on sending DNS requests to trigger routing loops. In IPv4, we probe all the routable prefixes retrieved from the RouteViews dataset [48], resulting in more than 3 billion individual IP addresses. In IPv6, however, the exhaustive scan of the routable space is not feasible. Instead, we scan more than 445 million hosts from the IPv6 Hitlist Service [17].

Figure 2 shows the measurement setup for the IPv4 scan. We run our experiments on top of the existing measurement infrastructure and use our custom scanner<sup>1</sup> capable of sending DNS packets in bulk [53]. Nevertheless, any other DNS scanner such as `zdns`,<sup>2</sup> would achieve the same goal. We set up an authoritative nameserver for `dnsdrakkarv4.com` domain name and all its subdomains. We encode the two following pieces of information in each queried domain: a random string (`xxxxxxx` in Figure 2) and the hexadecimally-encoded IPv4 address of the query target (`01020304` for 1.2.3.4). In IPv6, we encode the target IPv6 address as a network byte order 32-bit integer. The encoded address is used to attribute each domain name to the scanned destination address. As a result, all domain names uniquely identify each sent request. Importantly, we capture

<sup>1</sup> <https://github.com/mskwarek/myDig>

<sup>2</sup> <https://github.com/zmap/zdns>

**Table 1.** Repeating responses received on the scanner (October 2021).

Group	Response Count	Destination IP addresses	Destination ASNs	Average PAF	Maximum PAF
2 - 9 responses	15,511	15,488	1,733	2.1	9
10 - 254 responses	380	372	21	49.8	246
255 + responses	64	64	5	927,796	46,734,052

all the incoming requests on the authoritative nameserver and all the responses on the scanner.

We have run our scans from one vantage point. Although we plan to acquire more vantage points at different locations, we later show in Section 5 that the great majority of all the routing loops involve destination autonomous systems, so they can be triggered regardless of the measurement vantage point.

## 5 Scan Results

In this section, we first present the results of the latest Internet-wide IPv4 DNS scan (Section 5.1). We next run traceroute measurements towards the biggest amplifiers and identify routing loops (Section 5.2). Finally, we present the results of our two-year DNS measurement study in IPv4 and IPv6 (Section 5.3).

### 5.1 Internet Scan

We launched the latest Internet-wide IPv4 DNS scan in October 2021. In total, we sent more than 3 billion DNS A requests (one to each routable IP address) and received 7.6 million responses on the scanner. From each DNS response packet, we retrieve the following fields: the queried domain name (remember that each domain name is globally unique as it encodes the destination IP address to which we send the request), the source IP address of the response (can be the same as the destination IP address or different, in case the destination is a transparent forwarder [40]) and the DNS response code. We refer to each *response* as a three-tuple (*source\_IP\_address*, *domain\_name*, *response\_code*). Whenever we see a response tuple more than once, we refer to it as a *repeating response*.

Table 1 presents the results. We assign each repeating response to one of the three groups (first column) based on the number of times the response was received. We received 15,955 unique repeating responses in total. The first group (2 – 9 repeating responses) is the largest one, although the average amplification factor remains low (2.1 packets). Previous work analyzed the queries on root nameservers and found many repeating (with different query IDs) and identical (with same query IDs) requests [9, 59]. These were most probably results of configuration errors. Consequently, such repeating requests could produce repeating responses to our scanning host. As suggested by Bock *et al.* [5], responses sent more than 10 times are likely to be triggered by routing loops. If the TTL of the initial request is gradually decreased to 0, such a loop is finite. These responses

**Table 2.** Top 10 destination organizations (anonymized) in terms of triggered repeating responses.

Rank	Organization type	Country	Response count
1	Telecommunications Service Provider	PH	59,288,099
2	IT Services	GB	50,265
3	Internet Service Provider	IN	45,579
4	DNS services	CN	8,042
5	IT Services	US	5,390
6	Telecommunications Service Provider	CN	3,474
7	Internet Service Provider	CN	1,637
8	Telecommunications Service Provider	BR	956
9	Telecommunications Service Provider	IN	695
10	Telecommunications Service Provider	RU	624

belong to the second group (10 – 254 repeating responses). Note that in this case, the maximum count of received responses (254 responses) is an overestimation, as we would need to subtract from the maximum TTL (255 hops) the number of hops to reach the amplifier [5]. Finally, the third group (255+ repeating responses) contains the smallest number of response tuples, but the average PAF is very high (927,796 packets). The biggest amplifier seen during this scan triggered 46.7 million responses during 7 hours.

We use the CAIDA’s AS Rank dataset [7] to map autonomous system numbers (ASNs) to organization names and countries. All the destination autonomous systems originate from 133 countries, mostly from Brazil, India, and the USA. Table 2 presents the top 10 organizations (anonymized) in terms of the number of triggered repeating responses. The number one of the ranking (a Philippine telecommunications service provider) triggered many more responses to our scanning host than any other autonomous system.

We would expect that one DNS A request triggers repeating responses from the same source IP address and of the same DNS response type. In other words, one request triggers one repeating response tuple. Nevertheless, in groups 1 and 2, there are more repeating response tuples (second column of Table 1) than scanned destination IP addresses (third column of Table 1). The reason is that certain DNS requests triggered replies from different IP addresses. In particular, we found 15 destination IPs triggering repeating responses from 2 or more source addresses. Park *et al.* [42] have shown how a single request to the DNS forwarder was processed by 89 different recursive resolvers (as seen on the authoritative nameserver). If such a forwarder is transparent (i.e., it forwards the request without changing the source IP address field), the replies from different recursive resolvers will be returned to the original requester. Consequently, we could cumulate PAFs from all responses triggered by a single request.

The received DNS responses are of five following types (as defined in RFC-1035 [35]): `NOERROR` (13,797 responses), `SERVFAIL` (1,684 responses), `REFUSED` (430 responses), `NXDOMAIN` (41 responses) and `NOTIMP` (3 responses). We take a closer look at 345 `NOERROR` responses from groups 2 and 3. Although this response code signals that the request was completed successfully, the answer section of the DNS packet may have been manipulated. Surprisingly, 76% of

these responses did not contain the A record in the answer at all. As for the remaining non-empty responses, we did not detect any manipulation.

If the majority of NOERROR responses are empty, we raise the question of whether the authoritative nameserver for our test domain `dnsdrakkarv4.com` experienced any significant load from repeating requests. Specifically, we have analyzed 444 repeating responses from groups 2 and 3. As expected, the great majority of domain names (350 domains) were never queried on the authoritative nameserver, which suggests that attackers can safely reuse existing domain names in their queries without domain name operators noticing any abnormal activity.


In the attempt to characterize the devices responsible for response packet amplification, we made an assumption that we might have been dealing with national censors' middleboxes, as suggested by Bock *et al.* [5]. Censored Planet initiative [10] is constantly measuring the presence of censorship worldwide. More specifically, their Hyperquack [56] project infers application-layer blocking. We checked whether the Hyperquack data contains measurements towards destinations that triggered response floods but we did not find any overlap between the two datasets. We next referred to Tracebox [13] – a middlebox detection software that relies on ICMP `time-exceeded` replies to check whether the originally sent packet was modified and in which way. We run tracebox towards all the destination IP addresses from groups two and three. We notice modifications (such as unexpected source/destination addresses or ports, checksum, etc.) on the way to each measured host. However, when compared to a random sample of routable IP addresses (that did not trigger any amplification), there are no specific packet modifications that would distinguish the amplifier group from non-amplifiers. Therefore, identifying and characterizing those devices that trigger response floods remains an open question.

## 5.2 Running Traceroute


The responses from the second and third groups (see Table 1 rows 2 and 3) are likely to be caused by routing loops. One could use the traceroute [29] utility to track the path a packet takes from the source to the destination and check for the presence of loops. Augustin *et al.* [2] indicated, however, that traceroute does not capture the complete view of the network, often showing anomalies (such as loops) when there is router load balancing in place. To address this limitation, we use Multilevel MDA-Lite Paris Traceroute [57]. This tool relies on the new MDA-Lite algorithm to avoid inferring false links and to give a more accurate view of the path between the measurement server and the destination.

We trace the path to all 435 *unique* destination IP addresses in groups two and three (see Table 1) immediately after the end of the scan. Notice that certain destination IP addresses form more than one response tuple and, consequently, may appear in multiple groups, which is actually the case for one destination host that belongs to both groups two and three. The great majority of traceroutes (67%) did not reach measured destinations, even though they triggered repeating DNS responses. We found 115 unique loops towards 392 tested hosts.





**Fig. 3.** Highest amplification factors per individual IPv4 DNS scan. The three highest PAFs (note the logarithmic scale) are highlighted in red.



**Fig. 4.** Highest amplification factors per individual IPv6 DNS scan.

Importantly, Nawrocki *et al.* [39] have shown that roughly 90% of all the DNS DDoS events captured at the IXP used up to 100 amplifiers. Therefore, the discovered 115 loops are sufficient to mount real-world attacks. As for the remaining 43 destination IP addresses, we consider that packets may have encountered transient loops during the scan, which disappeared at the time of the traceroute measurements.

Traceroute loop lengths vary greatly and involve up to 38 interface IP addresses, but most often 2 (39 loops). Interestingly, 6 loops involved reserved IP addresses from private [36] and shared address ranges [58].

Overall, IP addresses involved in 115 routing loops originated from 35 autonomous systems. As for the location, 102 loops involved destination autonomous systems. Consequently, the great majority of all the routing loops could potentially be triggered from different vantage points.

### 5.3 Longitudinal Analysis

To test whether the threat of response floods is constantly present, we analyzed the results of regular DNS scans that we have been performing since February 2020 (22 IPv4 and 22 IPv6 scans). For each scan, we first identify the response returned the maximum number of times. We plot the highest packet amplification factors for IPv4 in Figure 3 (note that the y-axis is on a logarithmic scale). There are apparent outliers (highlighted in red): three scans generated 28, 69, and 655 million responses maximum, all triggered by sending one request to hosts in three different autonomous systems. Overall, these biggest response floods lasted between 7 seconds and 39 hours. However, occasionally resent A requests could keep restarting these loops.

We take a closer look at the maximum PAF ever observed. One query to a host from an autonomous system in the Philippines generated 655 million

**Table 3.** Repeating responses received by the scanner between February 2020 and December 2021 (IPv4 and IPv6 combined).

Group	Response Count	Destination IP addresses	Destination ASNs	Average PAF	Maximum PAF
2 - 9 responses	938,606	690,988	21,855	2.3	9
10 - 254 responses	4,750	2,132	295	40.8	254
255 + responses	731	542	42	1,852,087	655,195,124

SERVFAIL responses sent during 2 days. It is the same autonomous system that triggered most of the repeated responses during our latest scan. We performed a traceroute measurement towards this destination and found a routing loop involving 9 hosts from two autonomous systems (including the destination AS), 21 hops from the scanner. Note that this particular traceroute was limited to 64 hops. Overall, the maximum PAF per IPv4 scan varies between 8,795 and 655 million, the average maximum is 35 million.

On the other hand, for IPv6, the revealed amplification factors are less impressive. Figure 4 presents PAF for IPv6 (note that the y-axis is now in linear scale). The maximum PAF is 60, thus there are no infinite routing loops (unless the looping packet was dropped early). The average maximum amplification factor is 22. The IPv6 results should be interpreted with caution due to the composition of the IPv6 hitlist [17]. It contains responsive IPv6 addresses, whereas one of the root causes of routing loops is sending packets to announced but not allocated IP space. We performed an additional scan of randomly sampled 50 million IPv6 addresses from each routable /40 IPv6 network but did not trigger any routing loop.

Table 3 presents the same results as Table 1 in Section 5.1, but this time aggregated over two years. Similar to our latest scan, the great majority of repeating requests were sent between 2 and 9 times. The average amplification factors remain similar between groups 1 and 2, but the largest revealed amplifier significantly increased the average PAF of group 3. Altogether, the scanner received nearly 1 million repeated responses corresponding to 1.4 billion packets during two years. The destination IP addresses are distributed among 21,804 unique autonomous systems. More than half of autonomous systems in groups two and three appeared during two or more scans. Consequently, the routing loops on the way to these networks were very likely persistent and required a manual fix.

## 6 Ethical Considerations and Disclosure

Research scans are widespread these days, allowing for quick and efficient discovery of all sorts of vulnerabilities and misconfigurations. Nevertheless, measurement studies require careful planning so that risks are minimized and benefits outweigh potential inconveniences [14]. As there is no mechanism to explicitly request permission to scan each IP address in advance, researchers developed

a set of guidelines [15] to inform network operators about the scanning nature and opt-out easily. We follow those guidelines and configure our domain name (and all the subdomains) to point to a web page explaining who we are and what we do. The provided contact email address can be used to opt-out from future scans. In addition, we do not consecutively scan all the hosts of a single network but randomize our input. We received one complaint during the scan and removed 1 autonomous system from the experiment, containing 32k IPv4 addresses.

Discovered routing loops raise a significant threat to networks containing them and those receiving the response flood. We have used the Registration Data Access Protocol (RDAP) [16, 41] protocol to find contact information for the IP addresses involved in routing loops and notify the corresponding network administrators. In our emails, we explain how we discovered the vulnerability and the potential consequences.

## 7 Related Work

The number of open DNS resolvers dropped substantially in recent years – from 17.8 million in 2015 [26] to around 2 million in 2021 [4, 25, 40, 52]. Yet, DNS has been heavily involved in reflection and amplification attacks [20, 21]. In their recent work, Nawrocki *et al.* [39] extensively analyzed the whole DNS amplification ecosystem, using data from honeypots, an Internet Exchange Point (IXP), active measurements, and Internet-wide scans. They have shown that DNS-based amplification attacks are even more present than previously thought. Alarming, the attackers do not yet fully exploit all the available amplification potential.

One approach to the detection of DNS amplifiers is to craft a single request that will produce a large response. MacFarland *et al.* [34] issued **A** and **ANY** requests for 363 million (domain name, authoritative nameserver IP address) pairs to identify amplified responses. They reached a 32.77 amplification factor for **ANY** type query with EDNS0 enabled. It was later shown that **ANY** responses for DNSSEC-signed domains can reach the amplification factor of 179 [46].

Another approach is to create one DNS request that will trigger a series of additional lookups. The DNS Unchained attack requests recursive resolvers to follow a long chain of **CNAME** resource records [6]. Even more destructive is a recently discovered **NXNSAttack**, which relies on bogus referrals that can overwhelm both recursive resolvers and authoritative nameservers [1].

More generally, the stateless nature of UDP allows many protocols, apart from DNS, to be used for reflection and amplification. Rossow [47] analyzed 14 popular UDP-based protocols with amplification factors between 3.8 (BitTorrent, NetBios) and 4,670 (NTP). The latter, NTP, is infamous for its high DDoS potential and is often seen in real-world attacks [20]. Czyz *et al.* [11] estimated that roughly 2.2 million NTP servers could be misused. Earlier, Kühner *et al.* [27] cooperated with CERTs, NOCs, clearinghouses, and other security organizations worldwide to improve the NTP amplifier landscape.

It was long believed that the three-way handshake prevents TCP from being abused in reflection attacks with spoofed requests. In practice, one can trigger remote servers to retransmit (up to 20 times) unacknowledged SYN/ACK segments before the handshake is completed [27]. Additionally, other types of TCP misconfigurations (such as repeating RST packets or the actual data being transmitted before the handshake is completed) result in an average amplification factor of 112 [28]. Finally, recent work has gone beyond the initial handshake and found how network middleboxes can be used to reflect and amplify TCP traffic towards victims [5].

The method presented in this paper does not require any complex setup or specifically crafted requests to amplify the response. We rely on a trivial UDP packet to trigger routing loops.

## 8 Conclusions and Future Work

In this paper, we have shown how a single DNS A request can generate a response packet flood. We have scanned all the routable IPv4 address space and found 15,909 end-hosts in 1,742 autonomous systems that triggered the repeating responses with the maximum packet amplification factor of 46.7 million. We have collected traceroute measurements towards the destinations that triggered most responses and found 115 routing loops. We have disclosed our findings to network operators. Overall, having analyzed two years of our DNS scans, we have found 18 query destinations that triggered more than one million responses. The historical data reveals that this phenomenon is not a one-time event. At any instant, an attacker can locate amplifiers with little effort, trigger them, and redirect the generated traffic to a victim.

We foresee three directions for future work. First, we plan to identify and further characterize those devices triggering response packet floods. Second, we intend to perform scans from geographically distributed vantage points. Although we have shown that the majority of loops involve destination autonomous systems, there may be more loops in transit. Finally, we will explore which other query types and protocols can be used to trigger routing loops as easily as DNS.

## Acknowledgements

The authors would like to thank Baptiste Jonglez, the reviewers and our shepherd for their valuable and constructive feedback. This work was partially supported by RIPE NCC, Carnot LSI, the Grenoble Alpes Cybersecurity Institute under contract ANR-15-IDEX-02, and by the DiNS project under contract ANR-19-CE25-0009-01.

## References

1. Afek, Y., Bremler-Barr, A., Shafir, L.: NXNSAttack: Recursive DNS Inefficiencies and Vulnerabilities. In: 29th USENIX Security Symposium (USENIX Security 20). pp. 631–648. USENIX Association (Aug 2020)

2. Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Avoiding Traceroute Anomalies with Paris Traceroute. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement. p. 153–158. IMC '06, Association for Computing Machinery, New York, NY, USA (2006)
3. AWS Shield: Threat Landscape Report – Q1 2020, [https://aws-shield-tlr.s3.amazonaws.com/2020-Q1\\_AWS\\_Shield\\_TLR.pdf](https://aws-shield-tlr.s3.amazonaws.com/2020-Q1_AWS_Shield_TLR.pdf)
4. Bayer, J., Nosyk, Y., Hureau, O., Fernandez, S., Paulovics, I., Duda, A., Korczyński, M.: Study on Domain Name System (DNS) Abuse Final Report. Tech. rep. (2022)
5. Bock, K., Alaraj, A., Fax, Y., Hurley, K., Wustrow, E., Levin, D.: Weaponizing Middleboxes for TCP Reflected Amplification. In: 30th USENIX Security Symposium (USENIX Security 21). pp. 3345–3361. USENIX Association (Aug 2021)
6. Bushart, J., Rossow, C.: DNS Unchained: Amplified Application-Layer DoS Attacks Against DNS Authoritatives?. In: Bailey, M., Holz, T., Stamatogiannakis, M., Ioannidis, S. (eds.) Research in Attacks, Intrusions, and Defenses. pp. 139–160. Springer International Publishing, Cham (2018)
7. CAIDA AS Rank. <http://as-rank.caida.org/>
8. CAIDA: The Spoofer Project, <https://www.caida.org/projects/spoofers/>
9. Castro, S., Wessels, D., Fomenkov, M., Claffy, K.: A Day at the Root of the Internet. SIGCOMM Comput. Commun. Rev. **38**(5), 41–46 (Sep 2008)
10. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory, <https://censoredplanet.org>
11. Czyz, J., Kallitsis, M., Gharaibeh, M., Papadopoulos, C., Bailey, M., Karir, M.: Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In: Proceedings of the 2014 Conference on Internet Measurement Conference. p. 435–448. IMC '14, Association for Computing Machinery, New York, NY, USA (2014)
12. Deccio, C., Hilton, A., Briggs, M., Avery, T., Richardson, R.: Behind Closed Doors: A Network Tale of Spoofing, Intrusion, and False DNS Security. In: Proceedings of the ACM Internet Measurement Conference. p. 65–77. IMC '20, Association for Computing Machinery, New York, NY, USA (2020)
13. Detal, G., Hesmans, B., Bonaventure, O., Vanaubel, Y., Donnet, B.: Revealing middlebox interference with tracebox. In: Proceedings of the 2013 Conference on Internet Measurement Conference. p. 1–8. IMC '13, Association for Computing Machinery, New York, NY, USA (2013)
14. Dittrich, D., Kenneally, E.: The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Tech. rep., U.S. Department of Homeland Security (August 2012)
15. Durumeric, Z., Wustrow, E., Halderman, J.A.: ZMap: Fast Internet-wide Scanning and Its Security Applications. In: USENIX Security Symposium (2013)
16. Gañán, C.: WHOIS sunset? A primer in Registration Data Access Protocol (RDAP) performance. In: Network Traffic Measurement and Analysis Conference, TMA. IFIP (2021)
17. Gasser, O., Scheitle, Q., Foremski, P., Lone, Q., Korczynski, M., Strowes, S.D., Hendriks, L., Carle, G.: Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In: Proceedings of the 2018 Internet Measurement Conference. ACM, New York, NY, USA (2018)
18. Google Cloud: Exponential growth in DDoS attack volumes, <https://cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks>

19. Hengartner, U., Moon, S., Mortier, R., Diot, C.: Detection and Analysis of Routing Loops in Packet Traces. In: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement. p. 107–112. IMW '02, Association for Computing Machinery, New York, NY, USA (2002)
20. Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., Dainotti, A.: Millions of Targets under Attack: A Macroscopic Characterization of the DoS Ecosystem. In: Proceedings of the Internet Measurement Conference. p. 100–113. IMC'17, Association for Computing Machinery, New York, NY, USA (2017)
21. Kopp, D., Dietzel, C., Hohlfeld, O.: DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks. In: Hohlfeld, O., Lutu, A., Levin, D. (eds.) Passive and Active Measurement. pp. 284–301. Springer International Publishing, Cham (2021)
22. Korczyński, M., Nosyk, Y.: Source Address Validation. In: Encyclopedia of Cryptography, Security and Privacy. pp. 1–5. Springer Berlin Heidelberg, Berlin, Heidelberg (2019). [https://doi.org/10.1007/978-3-642-27739-9\\_1626-1](https://doi.org/10.1007/978-3-642-27739-9_1626-1)
23. Korczyński, M., Nosyk, Y., Lone, Q., Skwarek, M., Jonglez, B., Duda, A.: Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers. In: Proceedings of the Applied Networking Research Workshop. p. 9–11. ANRW '20, Association for Computing Machinery, New York, NY, USA (2020)
24. Korczyński, M., Nosyk, Y., Lone, Q., Skwarek, M., Jonglez, B., Duda, A.: The Closed Resolver Project: Measuring the Deployment of Source Address Validation of Inbound Traffic (2020)
25. Korczyński, M., Nosyk, Y., Lone, Q., Skwarek, M., Jonglez, B., Duda, A.: Don't Forget to Lock the Front Door! Inferring the Deployment of Source Address Validation of Inbound Traffic. In: Passive and Active Measurement. Springer International Publishing (2020)
26. Kühner, M., Hupperich, T., Bushart, J., Rossow, C., Holz, T.: Going Wild: Large-Scale Classification of Open DNS Resolvers. In: Internet Measurement Conference. ACM (2015)
27. Kühner, M., Hupperich, T., Rossow, C., Holz, T.: Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In: 23rd USENIX Security Symposium (USENIX Security 14). pp. 111–125. USENIX Association, San Diego, CA (Aug 2014)
28. Kühner, M., Hupperich, T., Rossow, C., Holz, T.: Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks. In: 8th USENIX Workshop on Offensive Technologies (WOOT 14). USENIX Association, San Diego, CA (Aug 2014), <https://www.usenix.org/conference/woot14/workshop-program/presentation/kuhrer>
29. linux.die.net: `traceroute(8)` - Linux man page, <https://linux.die.net/man/8/traceroute>
30. Lone, Q., Frik, A., Luckie, M., Korczyński, M., van Eeten, M., Ganan, C.: Deployment of Source Address Validation by Network Operators: A Randomized Control Trial. In: Proceedings of the IEEE Security and Privacy (S&P) (2022)
31. Lone, Q., Korczyński, M., Gañán, C., van Eeten, M.: SAVING the Internet: Explaining the Adoption of Source Address Validation by Internet Service Providers. In: Workshop on the Economics of Information Security (2020)
32. Lone, Q., Luckie, M., Korczyński, M., van Eeten, M.: Using Loops Observed in Traceroute to Infer the Ability to Spoof. In: Kaafar, M.A., Uhlig, S., Amann, J. (eds.) Passive and Active Measurement. pp. 229–241. Springer International Publishing, Cham (2017)

33. Luckie, M., Beverly, R., Koga, R., Keys, K., Kroll, J., claffy, k.: Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In: Computer and Communications Security Conference (CCS). ACM (2019)
34. MacFarland, D.C., Shue, C.A., Kalafut, A.J.: Characterizing Optimal DNS Amplification Attacks and Effective Mitigation. In: Mirkovic, J., Liu, Y. (eds.) Passive and Active Measurement. pp. 15–27. Springer International Publishing, Cham (2015)
35. Mockapetris, P.: Domain Names - Implementation and Specification. RFC 1035 (Nov 1987), <https://rfc-editor.org/rfc/rfc1035.txt>
36. Moskowitz, R., Karrenberg, D., Rekhter, Y., Lear, E., de Groot, G.J.: Address Allocation for Private Internets. RFC 1918 (Feb 1996)
37. Moura, G.C.M., Castro, S., Heidemann, J.S., Hardaker, W.: Tsunami: exploiting misconfiguration and vulnerability to ddos DNS. In: IMC'21: ACM Internet Measurement Conference. pp. 398–418. ACM (2021)
38. Moura, G.C.M., de Oliveira Schmidt, R., Heidemann, J.S., de Vries, W.B., Müller, M., Wei, L., Hesselman, C.: Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In: Proceedings of the 2016 ACM on Internet Measurement Conference. pp. 255–270. ACM (2016)
39. 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 2021 Internet Measurement Conference. IMC '21, Association for Computing Machinery, New York, NY, USA (2021)
40. Nawrocki, M., Koch, M., Schmidt, T.C., Wählisch, M.: Transparent Forwarders: An Unnoticed Component of the Open DNS Infrastructure. In: Proceedings of CoNEXT '21. Association for Computing Machinery, New York, NY, USA (2021)
41. Newton, A., Hollenbeck, S.: Registration Data Access Protocol (RDAP) Query Format. RFC 7482 (Mar 2015)
42. Park, J., Jang, R., Mohaisen, M., Mohaisen, D.: A Large-Scale Behavioral Analysis of the Open DNS Resolvers on the Internet. IEEE/ACM Transactions on Networking pp. 1–14 (2021)
43. Paxson, V.: End-to-End Routing Behavior in the Internet. In: Conference Proceedings on Applications, Technologies, Architectures, and Protocols for Computer Communications. p. 25–38. SIGCOMM '96, Association for Computing Machinery, New York, NY, USA (1996)
44. Paxson, V.: An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks. ACM SIGCOMM Computer Communication Review **31**(3), 38–47 (Jul 2001)
45. Reuters: Russia's Yandex says it repelled biggest DDoS attack in history, <https://www.reuters.com/technology/russias-yandex-says-it-repelled-biggest-ddos-attack-history-2021-09-09/>
46. van Rijswijk-Deij, R., Sperotto, A., Pras, A.: DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study. In: Proceedings of the Fourteenth ACM Internet Measurement Conference, ACM IMC 2014. p. 449–460. IMC '14, Association for Computing Machinery, New York, NY, USA (2014)
47. Rossow, C.: Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In: In Proceedings of the 2014 Network and Distributed System Security Symposium, NDSS (2014)
48. University of Oregon Route Views Project, <http://www.routeviews.org/routeviews/>
49. Sasaki, T., Ganan, C., Yoshioka, K., Eeten, M., Matsumoto, T.: Pay the Piper: DDoS Mitigation Technique to Deter Financially-Motivated Attackers. IEICE Transactions on Communications (2019)

50. Senie, D., Ferguson, P.: Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing. RFC 2827 (May 2000), <https://rfc-editor.org/rfc/rfc2827.txt>
51. The Shadowserver Foundation, <https://www.shadowserver.org>
52. Shadowserver Foundation: Open Resolver Scanning Project, <https://scan.shadowserver.org/dns>, retrieved: December 2021
53. Skwarek, M., Korczyński, M., Mazurczyk, W., Duda, A.: Characterizing Vulnerability of DNS AXFR Transfers with Global-Scale Scanning. In: IEEE Security and Privacy Workshops (SPW) (2019)
54. Sridharan, A., Moon, S.B., Diot, C.: On the Correlation between Route Dynamics and Routing Loops. In: Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement. p. 285–294. IMC '03, Association for Computing Machinery, New York, NY, USA (2003)
55. van der Toorn, O., Krupp, J., Jonker, M., van Rijswijk-Deij, R., Rossow, C., Sperotto, A.: ANYway: Measuring the Amplification DDoS Potential of Domains. In: Proceedings of the International Conference on Network and Service Management (CNSM). pp. 500–508 (2021)
56. VanderSloot, B., McDonald, A., Scott, W., Halderman, J.A., Ensafi, R.: Quack: Scalable remote measurement of application-layer censorship. In: USENIX Security Symposium (2018)
57. Vermeulen, K., Strowes, S.D., Fourmaux, O., Friedman, T.: Multilevel MDA-Lite Paris Traceroute. In: Proceedings of the Internet Measurement Conference 2018. p. 29–42. IMC '18, Association for Computing Machinery, New York, NY, USA (2018)
58. Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., Azinger, M.: IANA-Reserved IPv4 Prefix for Shared Address Space. RFC 6598 (Apr 2012)
59. Wessels, D., Fomenkov, M.: Wow, That’s a lot of packets. In: Passive and Active Network Measurement Workshop (PAM) (2003-04)
60. Xia, J., Gao, L., Fei, T.: Flooding Attacks by Exploiting Persistent Forwarding Loops. In: Internet Measurement Conference 2005 (IMC 05). USENIX Association, Berkeley, CA (Oct 2005)
61. Yazdani, R., van Rijswijk-Deij, R., Jonker, M., Sperotto, A.: A Matter of Degree: Characterizing the Amplification Power of Open DNS Resolvers. In: Passive and Active Measurement Conference (PAM) (2022)
62. Yoachimik, O.: Cloudflare thwarts 17.2M rps DDoS attack — the largest ever reported, <https://blog.cloudflare.com/cloudflare-thwarts-17-2m-rps-ddos-attack-the-largest-ever-reported/>
63. Zhang, M., Zhang, C., Pai, V., Peterson, L., Wang, R.: PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services. In: 6th Symposium on Operating Systems Design & Implementation (OSDI 04). USENIX Association, San Francisco, CA (Dec 2004)
64. Zhang, S., Liu, Y., Pei, D.: A Measurement Study on BGP AS Path Looping (BAPL) Behavior. In: 2014 23rd International Conference on Computer Communication and Networks (ICCCN). pp. 1–7 (2014)
65. Zhang, Y., Mao, Z.M.: Effective Diagnosis of Routing Disruptions from End Systems. In: 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI 08). USENIX Association, San Francisco, CA (Apr 2008)