

Evaluating DNS Resiliency and Responsiveness with Truncation, Fragmentation & DoTCP Fallback

Pratyush Dikshit*, Mike Kosek[†], Nils Faulhaber[†], Jayasree Sengupta*, and Vaibhav Bajpai[‡]

*CISPA Helmholtz Center for Information Security, Germany

[pratyush.dikshit | jayasree.sengupta]@cispa.de

[†]Technical University of Munich, Germany [kosek | nils.faulhaber]@tum.de

[‡]Hasso Plattner Institute, Germany [vaibhav.bajpai@hpi.de]

Abstract—Since its introduction in 1987, the DNS has become one of the core components of the Internet. While it was designed to work with both TCP and UDP, DNS-over-UDP (DoUDP) has become the default option due to its low overhead. As new Resource Records were introduced, the sizes of DNS responses increased considerably. This expansion of the message body has led to truncation and IP fragmentation more often in recent years where large UDP responses make DNS an easy vector for amplifying denial-of-service attacks which can reduce the resiliency of DNS services. This paper investigates the resiliency, responsiveness, and usage of DoTCP and DoUDP over IPv4 and IPv6 for 10 widely used public DNS resolvers. The paper specifically measures the resiliency of the DNS infrastructure in the age of increasing DNS response sizes that lead to truncation and fragmentation. Our results offer key insights into the management of robust and reliable DNS network services. While DNS Flag Day 2020 recommends 1232 bytes of buffer sizes, we find out that 3/10 resolvers mainly announce very large EDNS(0) buffer sizes both from the edge as well as from the core, which potentially causes fragmentation. In reaction to large response sizes from authoritative name servers, we find that resolvers do not fall back to the usage of DoTCP in many cases, bearing the risk of fragmented responses. As the message sizes in the DNS are expected to grow further, this problem will become more urgent in the future. This paper demonstrates the key results (particularly as a consequence of the DNS Flag Day 2020) which may support network service providers make informed choices to better manage their critical DNS services.

Index Terms—DNS, DNS-over-TCP, DNS-over-UDP, Response Time, Failure Rate, EDNS(0)

I. INTRODUCTION

The Domain Name System (DNS), which is responsible for the resolution of hostnames to IP addresses, has become one of the most widely used components on the Internet. Hostnames (domain names) are organized in a tree structure that is hierarchically separated into *zones*. The resolution of domain names is realized by different components such as stub resolvers, recursive resolvers, and authoritative Name Servers (NSes). While authoritative NSes are responsible for the authoritative mapping of domains in a zone to their IP addresses, stub, and recursive resolvers cache and deliver such information from the NSes to the clients via a DNS request [1] (RFC 1034 [2]). DNS communication supports both major transport protocols on the Internet, namely the Transmission Control Protocol (TCP) (RFC 793 [3]) and the User Datagram Protocol (UDP) (RFC 768 [4]). Due to its comparably low overhead, UDP has become the default transport protocol for DNS. The




Fig. 1: 2527 RIPE Atlas Probes communicate the DNS requests with the edge (Probe and Public Resolvers) and with the core (authoritative NSes) using IPv4 and IPv6. Cached DNS responses are sent by the edge, while uncached DNS responses (2KB and 4KB) are sent by the core.

UDP message body is restricted to 512 bytes (RFC 1035 [5]). However, the increase in deployment of DNS Security (DNSSEC) and IPv6 (RFC 7766 [6]) has resulted in larger message sizes, thereby leading to two important developments in the protocol. Firstly, DNS-over-TCP (DoTCP) was declared to be mandatory for hosts (RFC 5966 [7]) as it enables a larger message body by default. Secondly, Extension Mechanisms for DNS (EDNS) were introduced to augment the capabilities of the DNS protocol in terms of message size expansion (RFC 2671 [8]). With the new EDNS capability, it was required that DNS replies would continue to provide responses as UDP datagrams even though the response was larger than 512 bytes. Stipovic *et al.* in [9] examines the level of compatibility for a number of public DNS servers for some popular Internet domains while exploring the behavior of some contemporary DNS implementations such as Microsoft Windows 2012, 2016 and 2019 as well as Linux-based BIND in regards to the EDNS. However, using too large UDP buffer sizes can cause IP fragmentation in certain networks, thereby reducing resiliency in DNS communication [10]. To avoid fragmentation, the DNS Flag Day, 2020¹, an association of DNS software maintainers and service providers, recommended the usage of a default buffer size of 1232 bytes. DoTCP is a useful measure against fragmentation and can increase DNS resiliency by allowing fallback options. Resolvers should also avoid frag-

¹<http://www.dnsflagday.net/2020/>

mentation by using the recommended default EDNS(0) buffer size of 1232 bytes. To this end, our paper puts forward three **goals**: *a*) to evaluate DoTCP support (both over IPv4 and IPv6) and its usage across several DNS resolvers, *b*) to analyze the responsiveness/ latency over DoTCP and DoUDP for IPv4 and IPv6, and *c*) to investigate which buffer sizes are currently used in DNS traffic around the globe.

In pursuit of these goals, we evaluate the behavior of the resolvers from two different vantage points to investigate from the edge and from the core of the Internet to represent the communication of the resolvers with DNS clients and authoritative name servers. Firstly, DoTCP adoption, responsiveness, and EDNS(0) configuration are analyzed from the edge where the interaction between recursive resolvers and DNS clients running on the RIPE Atlas probes is measured. To scope DNS requests to the *Edge* of the network, we perform DNS queries for a domain that is likely cached by all resolvers, unlike in previous studies [11]. Secondly, the interaction of recursive resolvers with authoritative NSes is further studied. To allow DNS requests to leave the edge and move into the *Core* of the network, we provision dedicated NSes for a custom-crafted domain whose resolution is requested from the DNS resolvers. Using this methodology (see §III), we study failure rates, response times, EDNS options, and usage of DoTCP and DoUDP, as well as the EDNS(0) configurations both from the edge and the core (except response times analysis) that gives detailed insights into the potential resiliency of DNS communication on the Internet as depicted in Fig. 1. We perform measurements over IPv4 and IPv6 [12]. Overall, more than 14M individual measurements from 2527 RIPE Atlas Probes have been analyzed, highlighting that most resolvers show similar resiliency for both DoTCP and DoUDP. Our main **findings** (see §IV) are –

Resiliency from the edge: We observe that DoTCP (4.01%) tends to fail less often than DoUDP (6.3%) requests over IPv4. Contrarily, in the case of IPv6, we find a higher failure rate over both transport protocols (DoTCP 10%, DoUDP 9.61%). The analysis of response times for Public and Probe resolvers confirms the pattern of approximately doubled median response times for DoTCP compared to DoUDP in both the IP versions. This result is similar to the ones observed by Callejo *et al.* [13] while measuring latency during DNS resolution. We also observe that several public DNS resolvers still lack adoption (< 3.5%) of 1232B from the DNS Flag Day recommendation.

Resiliency from the Core: We find that DoTCP requests over IPv4 exhibit failure rates of 9.09% on public resolvers against higher failure rate of 11.53% over IPv6. Surprisingly, we find that the RIPE Atlas measurements ended successfully even after receiving a response with the TC-bit set, indicating a lack of proper fallback to DoTCP in many probes. Moreover, communication between resolvers and the authoritative NSes utilizes an EDNS(0) buffer size of 512 bytes less preferably (IPv4 0.24%, IPv6 0.13%) compared to the buffer sizes advertised to the RIPE Atlas probes (IPv4 27.41%, IPv6 26.04%). All DNS resolvers use EDNS(0) in most of the cases (> 99.84%). We also see other DNS options such as Cookie (4.80% IPv4, 7.91% IPv6) and EDNS Client Subnet (ECS)

(1.81% IPv4, 1.49% IPv6) advertised by the public resolvers, while Google mostly uses ECS (14.24% IPv4, 12.53% IPv6).

DoTCP Usage Rates: We observe that when 2KB responses are received from the NSes, all resolvers that mainly use canonical (see §IV scenarios, use TCP in their last request for >95% of the cases. In situations where 4KB responses are received, we observe that almost all resolvers use TCP in the vast majority of measurements over both IP versions (>98%).

This paper builds upon our previous study [14]. In this paper, we have additionally added significant background information (see: §II) related to the monitoring and performance evaluation of DNS query-response over both TCP and UDP transport protocols. DNS response times can be a critical metric when using DoTCP fallback, we, therefore, conduct further measurements comparing DoTCP and DoUDP response times from the edge of the network (see: §IV-B2). Subsequently, when evaluating from the core, we present additional insights by including a detailed analysis using `traceroute` and DoUDP response time for public resolvers (see: §IV-B3). Additionally, we perform a deep dive by measuring the number of successful responses that do not contain any valid ANSWER sections for the DNS queries (see: §IV-C4). Notably, our investigation reveals instances where RIPE Atlas measurements have successfully terminated, despite receiving a response with the TC-bit set, thereby, indicating a lack of proper fallback to DoTCP across multiple probes. Towards the end, we discuss the implications of this research on network and service management in §V followed by discussing the limitations of our study and highlight future research directions in §VI, followed by a broad discussion in §VII and concluding statements in §VIII.

II. BACKGROUND AND RELATED WORK

A. DNS Measurement

To measure DNS failure rates, DNS performance, and the buffer sizes used, several studies have been conducted in the last few years. Some of them are discussed here.

1) **Fragmentation:** With the increased message sizes, DNS queries can exceed the MTU of many networks. Moura *et al.* in [15] analyzes the fragmentation rates of DNS queries to the `.nl` top-level domain showing that less than 10k of 2.2B observed DNS responses by authoritative NSes are fragmented. Although fragmentation is in general fairly rare in DNS communication, the consequences can have negative effects on the resiliency and connectivity of Internet applications (RFC 8900 [16]). Herzberg and Shulman in [17] presented an attack allowing to spoofing of Resource Records (RRs) in fragmented DNS responses by which attackers can hijack domains or nameservers under certain conditions. Following a similar procedure, Shulman and Waidner in [18] showed the opportunity to predict the source port used by the client. Both of these approaches belong to the class of DNS cache poisoning attacks, one of the most common and dangerous on the DNS. Though, cache poisoning attacks are also possible when DNS messages are not fragmented [19][20][21]. The aforementioned studies however show the additional security risk caused by fragmented responses. This potentially exposes

the DNS user to several other types of attacks. Koolhaas *et al.* in [10] analyzed the behavior of different EDNS(0) buffer sizes in 2020. It was shown that the likelihood of a failing DNS query increases with growing buffer sizes. For a size of 1500 bytes, the default MTU of Ethernet which causes fragmentation of most of the DNS messages of the DNS queries to stub resolvers failed for 18.92% over IPv4, with 26.16% over IPv6 (RFC 2464 [22]). As countermeasures, in 2017, Cao *et al.* in [23] presented an "Encoding scheme against DNS Cache Poisoning Attacks". Berger *et al.* in [20] presented a way of detecting DNS cache poisoning attacks in 2019. Even though it was later shown that DNS cache poisoning attacks are also possible against DoTCP by Dai *et al.* in [24], this emphasizes its importance as a fallback option to the usage of DoUDP. Herzberg and Shulman in [18] recommend keeping the indicated buffer size less or equal to 1500 bytes. As a consequence, Weaver *et al.* summarize a list of recommendations to stakeholders in the DNS ecosystem in [25]. These include the proposition for stub resolvers as well as authoritative nameservers to stick to buffer sizes of 1400B and below. The study conducted in [10] in 2020 yielded detailed recommendations for the EDNS(0) buffer size configuration of authoritative name servers and stub resolvers dependent on the used IP version and network type. The recommendations were adopted at the DNS Flag Day, 2020 claiming that "defaults in the DNS software should reflect the minimum safe size which is 1232 bytes". The aforementioned aspects emphasize the need for DNS resolvers to adopt the buffer size recommendations as fast as possible. Some other encrypted DNS protocol implementations such as DNS-over-TLS (RFC 7858 [26]) and DNS-over-HTTPS (RFC 8484 [27]) also counter the problems of fragmentation as TCP used as transport protocol [28]. They are however not yet widely enough adopted to obsolete standard DNS implementations [29][30][31]. To investigate the progress of DNS resolvers in implementing the new standards, measurements analyzing the buffer sizes used by DNS resolvers are therefore performed from different standpoints. As DoTCP support is another important requirement for DNS resolvers to avoid truncation and fragmentation, DNS failure rates over TCP and UDP are analyzed in this paper. Additionally, the DoTCP-fallback behavior of the resolvers is studied to see in which cases they make use of TCP. As furthermore response time is the main disadvantage when using DoTCP instead of DoUDP, we are also interested in comparing the two implementations with regard to this aspect.

2) **Response Times and Failure Rates:** The first large-scale study on DNS performance and failure rates was performed by Danzig *et al.* in 1992 resulting in several important recommendations to reduce DNS traffic and latency [32]. Later, DNS Performance and the Effectiveness of Caching analyzed DNS traffic of the MIT Laboratory for Computer Science and the KAIST (Korea Advanced Institute of Science and Technology) over several months and stated a failure rate of 36% (23% timeouts, 13% other errors) as explained by Jung *et al.* in [33]. Ager *et al.* [34] compared DNS response times between local DNS resolvers of ISPs and public resolvers like Google, PublicDNS, and OpenDNS. It was found that

local resolvers generally outperformed public resolvers, but Google and OpenDNS showed faster responses in certain cases due to ISP caching issues. Several other measurements have been undertaken to observe DNS performance over IPv4 from different standpoints [35][36][37]. Additionally, Doan *et al.* in [38] observed that the public resolvers answered faster over IPv6 than over IPv4. Callejo *et al.* in [13] study the impact on performance in terms of DNS resolution delays for both cached and non-cached domain names. For this, they discuss EDNS options announced to the ANSes. However, their study is limited to the EDNS client subnet (ECS) option only where they highlight Cloudflare, Quad9, and DNS.SB, all of which do not support ECS, except Google. They use a public resolver over DoH instead of using a local resolver over Do53. The study shows that the median resolution time increases by 17 ms when using Cloudflare/DoH, 40 ms for Quad9, 70 ms for Google/DoH, and 170 ms for DNS.SB. In 2022, Moura *et al.* [15] investigated the fallback capabilities of DNS resolvers to utilize DoTCP by manipulating the TC-bit in responses from a controlled authoritative name server. They introduced the distinction between canonical (UDP request followed by TCP request) and non-canonical scenarios. Evaluating the order of incoming requests, it was concluded that an estimated 2.7% (optimistic estimation) to 4.8% (pessimistic estimation) of the examined resolvers were incapable of falling back to DoTCP usage. In the same year, Kosek *et al.* [11] conducted the first study comparing DNS response times and failure rates based on the underlying transport protocols using RIPE Atlas for ten public resolvers and probe resolvers. 8% of the queries over UDP and TCP failed, with a very high DoTCP failure rate of 75.0% for probe resolvers. The response times of DoTCP were generally higher than that of DoUDP with large differences. The queried domains were unique and thereby uncached by all of the participating resolvers. To study the Internet measurement platforms for large-scale measurements, Sermpezis *et al.* in [39] showed that the bias of Internet measurement platforms differs significantly by dimension where RIPE Atlas is substantially less biased than RIPE RIS and RouteViews. They study the biases in RIPE Atlas, RIPE RIS, and RouteViews to compare the set of ASes hosting RIPE Atlas probes (both IPv4 and IPv6) and conclude that the set of networks hosting IPv6 probes is slightly more biased than networks hosting IPv4 probes in most dimensions. To get a preferably broad and unbiased comparison of the different DNS resolvers over the particular underlying protocols, we consider only probe-resolver pairs for successful responses and perform DNS queries over RIPE Atlas probes to both the domains, which is very likely cached on each resolver (*google.com*), and uncached ones. Additionally, we make sure that the uncached domains are administered by an authoritative name server under our control. Using a cached domain allows a detailed estimation of the latencies in the direct communication between resolvers and client software without the additional time needed for a recursive lookup. The DNS queries for uncached domains force recursive resolvers to forward them to our server. Observing the incoming requests offers the opportunity to analyze the communication between recursive resolvers and name servers, for example, the usage

TABLE I: *The public DNS resolvers under investigation, IPv4 and IPv6 represent the resolvers' anycast addresses*

Name	IPv4	IPv6
Cloudflare Public DNS	1.1.1.1	2606:4700:4700::1111
Google Public DNS	8.8.8.8	2001:4860:4860::8888
CleanBrowsing	185.228.168.9	2a0d:2a00:1::
OpenDNS	208.67.222.222	2620:119:35::35
OpenNIC	185.121.177.177	2a05:dfc7:5::53
Quad9	9.9.9.9	2620:fe::fe
Comodo Secure DNS	8.26.56.26	-
UncensoredDNS	91.239.100.100	2001:67c:28a4::
Neustar UltraDNS Public	64.6.64.6	2620:74:1b::1:1
Yandex.DNS	77.88.8.8	2a02:6b8::feed:ff

of DoTCP on the same path, in detail.

3) **EDNS Options:** Van den Broek *et al.* [40] analyzed more than 8 million DNS queries to an authoritative name server in 2014. Around 75% of the queries used EDNS(0). Additionally, it was observed that 36% used a UDP buffer size higher than 1232 bytes likely causing fragmentation. Measurements after the DNS Flag Day 2020 recommendations show that DNS resolvers still seem to lack adoption to the default buffer size of 1232 bytes. Based on the analysis of 164 billion queries to authoritative name servers, Moura *et al.* [15] stated that many resolvers "announce either small (512 bytes) or large (4096 bytes) EDNS(0) buffer sizes, both leading to more truncation and increasing the chances of fragmentation/packets being lost on the network". As the DNS Flag Day 2020 recommendations have not been out in the community for a very long time, a regular examination of the adoption rates of DNS software is reasonable and necessary.

III. METHODOLOGY

To extend the previous studies by Kosek *et al.* [11], we utilize identical target resolvers for our measurements. We query the ten public resolvers listed in Table I, along with the configured probe resolvers. It is worth mentioning that at the time of the initial measurements, Comodo Secure DNS did not have a known IPv6 address, resulting in measurements conducted solely over IPv4.

A. Probe Selection

This paper employs the RIPE Atlas measurement network to conduct the measurements. To avoid potential load issues occurring in the first two probe versions [41][42], we choose only probes of version 3 or 4 that are hosted with a hometag [43]. The chosen probes must support IPv4, IPv6, or both. The scan of RIPE Atlas probes was conducted on 20th December 2021 which produced the list of available probes. The performance measurements from the edge and the core were then conducted from 21st to 28th December 2021 and 19th to 25th March 2022 respectively. Based on the obtained results, the analysis was performed in 2022. All the measurements were conducted in a single stretch over a week. The scan revealed the availability of 2527 probes. Out of these, 1137 probes are IPv6 capable, while still supporting IPv4. These probes are distributed across 671 different Autonomous Systems (ASs) with varying densities

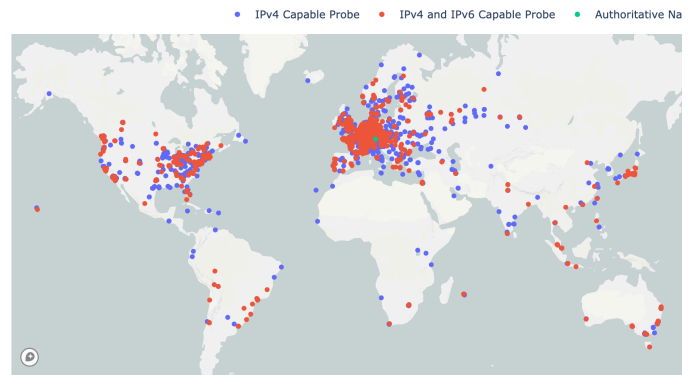


Fig. 2: *The global distribution of the RIPE Atlas probes participating in the experiments and the authoritative name servers developed for the measurements from the core.*

in different regions: 70% in Europe, 18% in North America, 6% in Asia, 3% in Oceania, 1% in Africa, and 1% in South America. It is to be noted that our measurements encounter some bias due to the uneven placement of the vantage points on the RIPE Atlas. The distribution of the probes among ASes participate in our measurements, 27% of the IPv4-capable ones are located in the top ten ASes (even 38% of the IPv6-capable probes). This yields a bias towards the most used ASes.

Before commencing the actual measurement series, which includes analyzing DNS resolvers from the edge and the core (as depicted in Fig. 1), examining DoTCP usage, EDNS(0) configuration, and DoTCP fallback, we evaluated 4343 probe resolvers associated with the 2443 participating probes, resulting in an average of 1.78 resolvers per probe.

B. From the edge

We programmatically configured RIPE Atlas measurements targeting DNS resolvers listed in Table I, covering both IPv4 and IPv6 and using TCP and UDP transport protocols. Treating the resolvers as black boxes, we focused on direct communication between DNS clients and recursive resolvers, without examining specific transport protocols or buffer sizes for authoritative name server communication. Measurements from the edge involved requesting "A" records for the widely-used domain "google.com," aiming to minimize recursive resolution and emphasize client-resolver communication. Response times collected by RIPE Atlas probes correspond to simple UDP/TCP request durations. RIPE Atlas probes collected response times, error messages, and UDP buffer sizes, retrieved via the RIPE Atlas measurement API. Simultaneous Traceroute measurements were conducted to validate response times and gain insights into resolver distribution across the Internet, anticipating Traceroute RTT alignment with DoUDP request response times for cached domains.

C. From the core

This evaluation allows us to analyze resolver behavior when interacting with authoritative NS (see Fig. 2). In this experiment, we use uncached domain names controlled by authoritative NS under our supervision. We analyze the resolvers' DNS configuration using two customized authoritative

TABLE II: *Probe distribution over Autonomous Systems along with the percentage of IPv4/IPv6 capable probes.*

ASes	IPv4 capable	% of probes	IPv6 capable	% of probes
DTAG	137	5.42%	109	9.59%
COMCAST	100	3.96%	61	5.37%
VODANET	96	3.80%	48	4.22%
PROXAD	89	3.52%	76	6.68%
Orange S.A.	75	2.97%	58	5.10%
AT&T	51	2.02%	28	2.46%
UUNET	37	1.46%	7	0.62%
TNF-AS	34	1.35%	15	1.32%
LDCOMNET	33	1.31%	10	0.88%
KPN	30	1.19%	17	1.50%

NSes. These NSes encode incoming DNS requests, including transport protocol and requester IP address, for later analysis. By observing the EDNS section of requests reaching the authoritative NS, we gain insights into the resolvers' EDNS configuration and potential usage of options like *Cookie* or *Client Subnet*.

1) **DoTCP Usage and EDNS(0) Configuration:** In the Core analysis, we utilize the RIPE Atlas network to evaluate the transport protocols and EDNS configuration employed by DNS resolvers. Our measurements focus on a target domain managed by our NSes to ensure uncached responses. To achieve uniqueness, each DNS request is modified with the probe ID and timestamp. By examining the IP addresses of the resolvers, we gain insights into their distribution across continents and AS networks. Additionally, observing the transport protocols used enables us to gather information on DoTCP usage in the Core.

2) **DoTCP Fallback:** The Core measurement focuses on observing the DoTCP fallback behavior of public DNS resolvers. Large responses, consisting of 72 AAAA records (>2KB response) for one server and 145 AAAA records (>4KB response) for the other, are returned by the authoritative NSes. By including different RR types (A, AAAA, and TXT) in each measurement, we aim to investigate the resolvers' reaction to both response sizes simultaneously. Given that the previous experiment revealed resolvers requesting both NSes equally, approximately 50% of the requests are expected to receive 4KB and 2KB responses respectively. As large responses cannot be handled by UDP due to fragmentation issues, resolvers are anticipated to fallback to using DoTCP. The analysis focuses on whether the resolvers continue to utilize UDP or switch to DoTCP, providing insights into potential resiliency risks. Multiple requests from resolvers to the authoritative NS are expected, such as one over UDP followed by a fallback request over TCP. To accurately map incoming requests to the RIPE Atlas measurements, the domains queried by each probe are made unique using the aforementioned technique of prepending probe-specific information.

Reproducibility – In order to enable the reproduction of our findings, we make the raw data [44] of our measurements, as well as the developed tools, analysis scripts and supplementary files publicly available [45].

IV. RESULTS

We evaluate the results of the measurement from the edge concerning failure rates, response times, and EDNS(0) buffer sizes. Afterward, we analyze EDNS(0) configuration and DoTCP fallback from the core.

A. Probes

A comprehensive examination of all RIPE Atlas probes reveals the presence of 2527 probes possessing the desired attributes. All probes exhibit compatibility with IPv4, while 1137 probes exhibit the additional capability of conducting measurements over IPv6. The geographic distribution of the probes, as well as the locations of the authoritative name servers, is visually depicted in Fig. 2. Notably, a concentrated density of probes is observed in North America and Europe, which serves as the primary origin for the RIPE Atlas community. Specifically, Europe accounts for 70% of the probes, followed by North America with 18%, Asia with 6%, Oceania with 3%, Africa with 1%, and South America with 1%. The distribution of these probes among the Autonomous Systems is detailed in Table II, highlighting the ten Autonomous Systems housing the majority of probes. The probes from Comcast, AT&T, and UUNET are exclusively situated in North America, while the remaining Autonomous Systems are primarily distributed in Europe. Furthermore, it should be noted that certain Autonomous Systems mentioned in Table II possess only a small number of IPv6-capable probes. The mapping between Autonomous System numbers and their respective names is obtained from IPtoASN¹.

Probe Resolvers: Before conducting the actual measurement, preliminary test measurements are performed to gather address information regarding the locally configured resolvers on each probe. It is important to note that the resolvers can possess either publicly accessible IP addresses or private ones. For instance, a probe may have Google Public DNS configured, in addition to a DNS resolver exclusively accessible through its local network. Among the probes, 41.93% have public IP addresses. Out of all the registered resolvers, 97.03% are associated with their corresponding IPv4 addresses, while the remaining 2.97% are linked to their IPv6 addresses. Additionally, it is crucial to acknowledge that during a RIPE Atlas measurement employing the utilized probe resolver parameter set, all probe resolvers are requested to resolve the relevant domain name. Consequently, both IPv4 and IPv6 DNS resolvers contribute to the measurement, regardless of the IP version mandated by the RIPE user.

B. Evaluation from the edge

This section analyzes failure rates of transport protocols at the edge to assess DNS resilience. Response times of resolvers are compared and the performance of public resolvers is evaluated using DNS response times and traceroute round-trip times. The adoption of DNS Flag Day recommendations by individual resolvers is analyzed through EDNS(0) buffer size announced to RIPE probes.

¹<https://iptoasn.com>

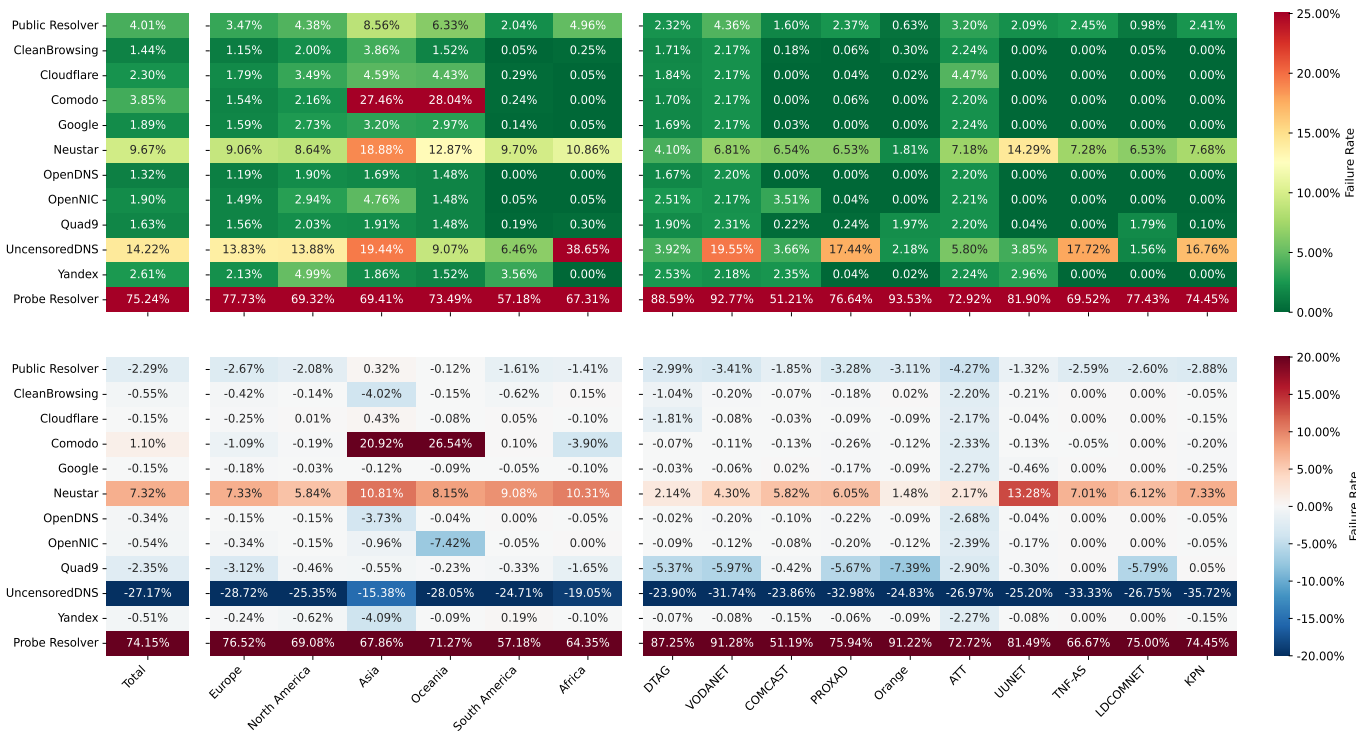


Fig. 3: Failure rates observed from the edge over IPv4. The upper part represents the DoTCP failure rates of all resolvers in total and per continent and AS. The lower part reflects the difference between the DoTCP and the DoUDP failure rates for a particular pairing (a negative value hence indicates a higher DoUDP failure rate). 'Public Resolver' summarizes the observations of all resolvers that are not probe resolvers.

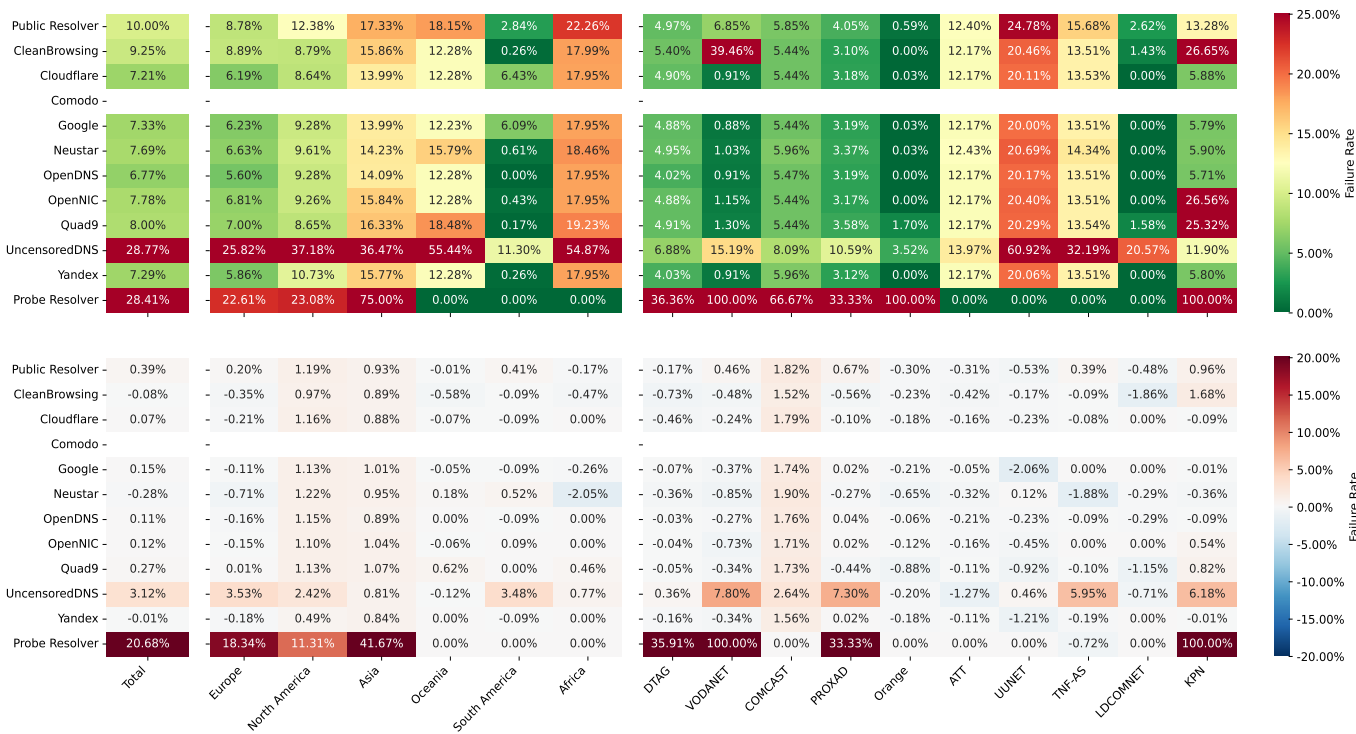


Fig. 4: Failure rates observed from the edge over IPv6. The upper part presents the failure rates over DoTCP, the lower one the difference between DoTCP and DoUDP failure rates. White cells indicate that there is no data for the given pairing.

TABLE III: EDNS(0) Buffer Sizes announced to the RIPE probes by the resolvers observed from the edge. CB: Clean-Browsing; U-DNS: UncensoredDNS

		512	1232	4096	none	other
CB	IPv4	97.04%	0.63%	1.46%	0.57%	0.30%
	IPv6	99.41%	0.11%	0.48%	0.01%	0.00%
Cloudflare	IPv4	0.20%	97.43%	1.45%	0.53%	0.40%
	IPv6	0.11%	99.44%	0.44%	0.01%	0.00%
Comodo	IPv4	0.18%	0.64%	98.30%	0.57%	0.30%
	IPv6	-	-	-	-	-
Google	IPv4	96.82%	0.78%	1.47%	0.58%	0.34%
	IPv6	99.22%	0.10%	0.67%	0.00%	0.01%
Neustar	IPv4	0.18%	0.64%	98.32%	0.56%	0.30%
	IPv6	0.10%	0.10%	99.79%	0.00%	0.00%
OpenDNS	IPv4	0.18%	0.63%	98.20%	0.57%	0.43%
	IPv6	0.10%	0.11%	99.79%	0.00%	0.00%
OpenNIC	IPv4	0.18%	97.53%	1.42%	0.56%	0.30%
	IPv6	0.11%	99.43%	0.47%	0.00%	0.00%
Quad9	IPv4	19.15%	55.47%	1.48%	23.55%	0.35%
	IPv6	20.98%	62.09%	0.47%	16.46%	0.00%
U-DNS	IPv4	0.30%	95.87%	2.39%	0.96%	0.49%
	IPv6	0.13%	99.29%	0.57%	0.01%	0.00%
Yandex	IPv4	0.19%	0.64%	98.04%	0.75%	0.39%
	IPv6	0.11%	0.10%	99.79%	0.00%	0.00%
Overall	IPv4	24.97%	36.12%	35.30%	3.24%	0.36%
	IPv6	24.86%	38.81%	34.46%	1.87%	0.00%

1) **Failure Rates:** Based on Kosek *et al.*'s research in [11], failed measurements are defined as those with no DNS response at the probe. In IPv4, public resolvers show lower failure rates for DoTCP (4.01%) compared to DoUDP (6.3%), indicating higher resiliency of DoTCP (Fig. 3). However, probe resolvers present a different scenario with DoTCP surpassing DoUDP by 74.15%. DoUDP failures are solely due to Timeouts (5000ms). And public resolvers' DoTCP failures are also primarily caused by Timeouts (42.75%), READ-ERROR (33.91%), CONNECT-ERROR (23.24%), and TCP-READ (0.09%). Bad address (99.17%) is the main cause of DoTCP failures for probe resolvers. Overall, probe resolvers exhibit significantly higher DoTCP failure rates across continents. In case of IPv6, for public resolvers, we find lower resiliency over both transport protocols (DoTCP 10%, DoUDP 9.61%). Most public resolvers exhibit failure rates between 6.77% and 9.25%, see Fig. 4. Uncensored DNS shows by far the worst DoTCP and DoUDP resiliency.

To analyze the adoption of the DNS Flag Day 2020 recommendations by the public resolvers from the edge, we evaluate the EDNS(0) buffer sizes which the individual resolvers announce to the RIPE Atlas probes. Table III summarizes the buffer sizes that have been observed in the UDP measurements. As for all resolvers except Quad9 the difference in the percentages of the announced buffer sizes between IPv4 and IPv6 are fairly low ($\leq 3.5\%$). The buffer sizes advertised by Cloudflare, OpenNIC, UncensoredDNS, and Quad9 (55.47% IPv4, 62.09% IPv6) conform to the DNS Flag Day 2020 recommendation of a default buffer size of 1232B in most cases. Neustar, Comodo, OpenDNS and Yandex mainly use 4096 bytes. In 23.55% of the Quad9 DNS responses over IPv4, EDNS(0) is not used at all leaving clients to the default DoUDP message size limit of 512 bytes. This first view from the edge shows that several public DNS resolvers still lack adoption to the DNS Flag Day 2020 recommendations. To see whether this also holds for the communication with authoritative NSes, we conducted another experiment from the core.

Takeaway: The probe resolvers demonstrate notably elevated DoTCP failure rates on a global scale. Specifically, in the context of IPv6 and public resolvers, a diminished resiliency is observed across both transport protocols, with DoTCP and DoUDP rates at 10% and 9.61%, respectively. The majority of public resolvers experience failure rates ranging from 6.77% to 9.25%.

2) **Response Times:** RIPE Atlas employs a measurement methodology to assess the response time (RT) of DNS requests. It measures the duration from the initiation of the measurement until a valid DNS response is received at the probe. Since many unicast resolvers can only be reached within their own local networks, and RIPE Atlas doesn't permit Traceroute measurements within such networks for security reasons, we were unable to gather Round-Trip Times (RTTs) for most unicast resolvers. This may seem to create unfairness in the measurement, but to counter its effects and ensure a fair comparison of response times, only probe-resolver pairs with successful responses over both TCP and UDP have been considered in our analysis.

DoUDP enables immediate transmission of requests to resolvers without the need to establish a TCP connection like DoTCP. Considering the cached nature of the "google.com" domain, DoUDP requests to resolvers with efficient cache management are expected to have response times equivalent to the probe-resolver round-trip time. In contrast, DoTCP measurements involve a three-way handshake, resulting in response times roughly twice as long as DoUDP. The depicted response times in Figures 5 and 6 are obtained through a two-tiered approach, calculating the median response time for each probe-resolver combination and presenting the median of all probes.

IPv4. The analysis of response times for Public and Probe resolvers confirms the expected pattern of approximately doubled median response times for DoTCP compared to DoUDP. Probe resolvers, due to their close physical proximity to the probing device, exhibit faster response times than Public resolvers. Among the Public DNS resolvers examined, Cloudflare, Google, and Quad9 demonstrate the lowest median response times for both transport protocols. Yandex shows relatively long response times for both DoTCP (104.5ms) and DoUDP (51.3ms). Neustar, on the other hand, exhibits the longest median DoTCP response time (1035.8ms) in this experiment, along with high DoTCP failure rates, suggesting inadequate implementation. When comparing response times by continent, Public resolvers generally respond fastest to DoTCP and DoUDP requests from European (DoTCP 52.7ms, DoUDP 24.1ms) and North American (DoTCP 54.9ms, DoUDP 27.1ms) probes.

The increased response times observed across various continent/resolver combinations for both transport protocols can be attributed to the sparser distribution of resolver Points-of-Presence (PoPs) in those continents. This is evident as DoTCP response times are consistently approximately twice as high as DoUDP response times, indicating longer Round-Trip Time (RTT) due to greater distances between probes and resolvers. As Sermpezis *et al.* [39] reveal a higher density of the

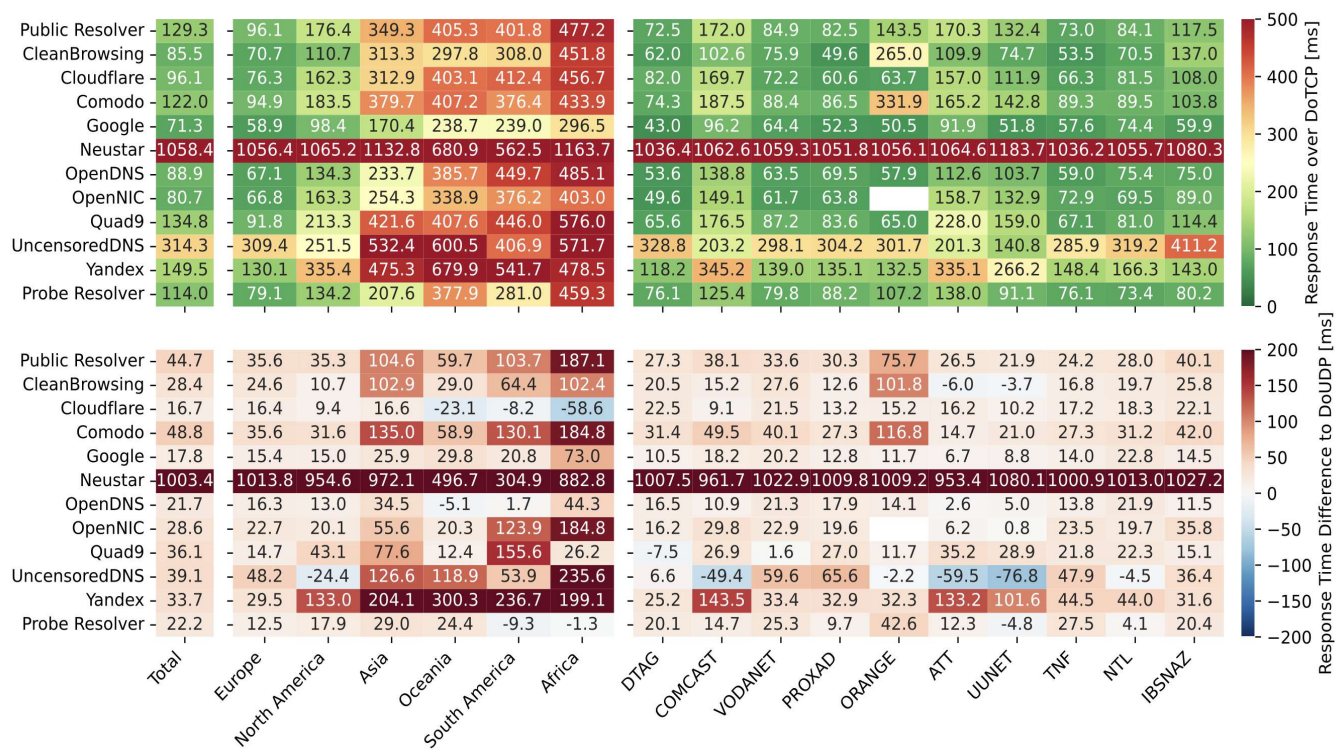


Fig. 5: Response times (RTs) observed from the edge over IPv4. The values represent the median RT over the medians of each probe. The upper part shows the results for DoTCP, the lower one the differences between DoTCP and DoUDP.

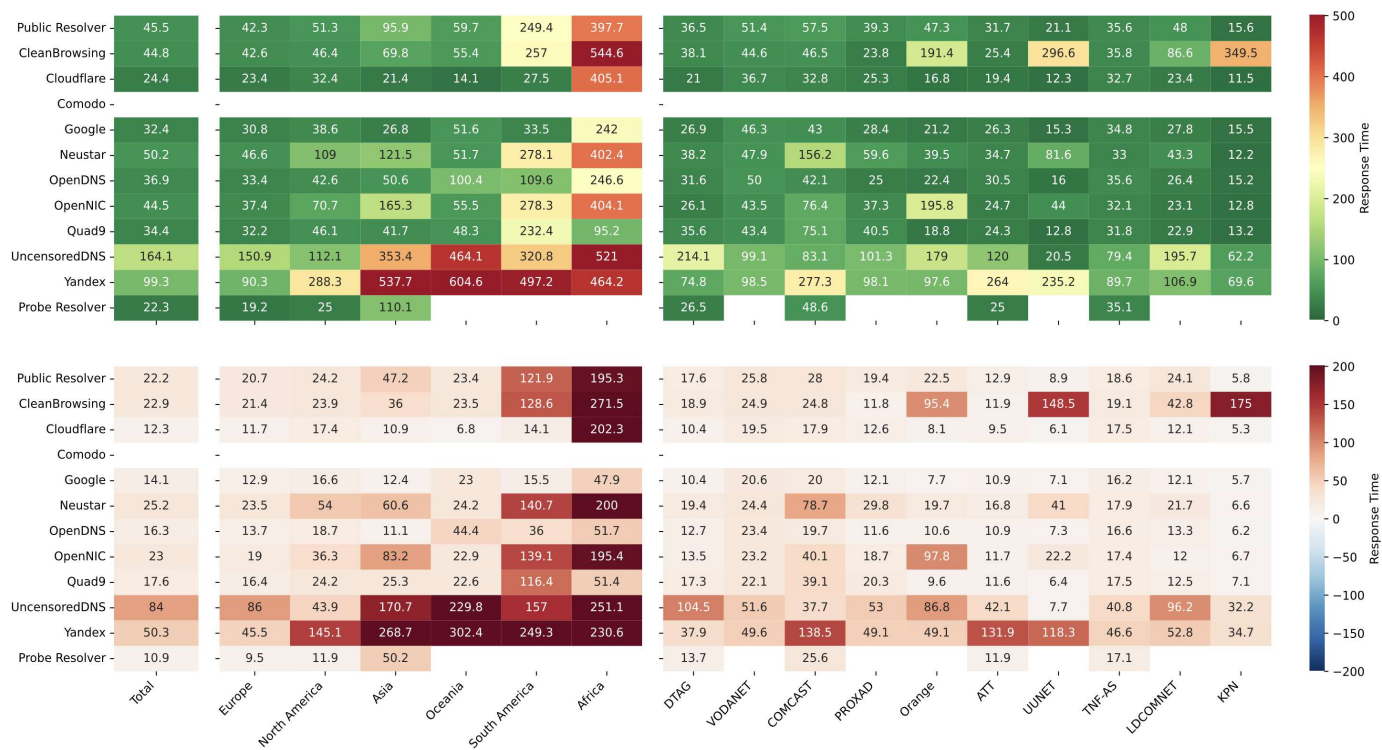


Fig. 6: Response times (RTs) observed from the edge over IPv6. The values represent the median RT over the medians of each probe. The upper part shows the results for DoTCP, and the lower one the differences between DoTCP and DoUDP. White cells indicate that there is no data for the given pairing.

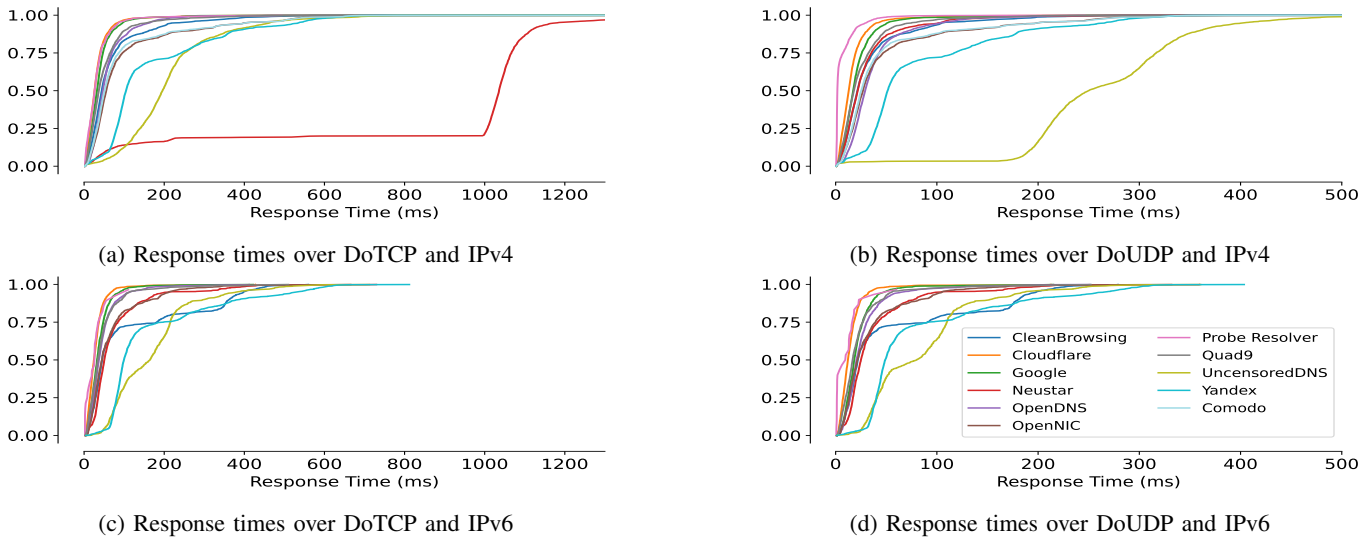


Fig. 7: Response times of each resolver over TCP/UDP and IPv4/IPv6 as CDF. The curves present an accumulation of the medians RTs of each probe.



Fig. 8: Traceroute round-trip-times (RTTs) and DNS response times (RTs) as a CDF. Again, the curves reflect an accumulation of probe-medians.

infrastructure in Europe, which is in imbalance with the spread of ASes around the world, i.e., there is location bias in RIPE Atlas probes, our strategy of considering only probe-resolver pairs for successful responses ensures a fair comparison, thereby minimizing the impact of bias in our results. We observe that probe resolvers demonstrate relatively low response times for both DoTCP (6.4ms - 32.2ms) and DoUDP (3.2ms - 15.3ms). Among public resolvers, Orange Autonomous System (AS) exhibits the highest overall median DoTCP response time (167.2ms) and also shows elevated DoTCP response times for CleanBrowsing, Comodo, and OpenNIC resolvers. Overall, most ASes, primarily operating in North America and Europe, do not exhibit significant anomalies in DoTCP and DoUDP response times. Yandex performs poorly for requests from North American ASes (223.9-290ms) but slightly better for European ASes (71.9-103.6ms). Conversely, UncensoredDNS displays higher DoTCP response times for European ASes (117.8-214.6ms) compared to North American ASes (36.6-145.9ms). Notably, UncensoredDNS consistently exhibits higher DoUDP response times than DoTCP, including a median DoUDP response time over seven times higher for requests from UUNET.

IPv6. The majority of Public resolvers show similar median response times for both DoTCP and DoUDP compared to their

IPv4 counterparts. However, there is a notable decrease in the median response time for DoTCP requests to Neustar (50.2ms) over IPv6. UncensoredDNS and Yandex still exhibit relatively high DoTCP response times. The improved performance of Neustar over DoTCP and UncensoredDNS over DoUDP contributes to lower overall median response times for public resolvers across both transport protocols. It is important to mention that probe resolvers are not considered in the IPv6 analysis. The analysis based on continents for Public resolvers reveals results similar to those observed in IPv4 measurements. Notably, the DoTCP response time for Africa is more than 70ms higher in IPv6 compared to IPv4. This can be attributed to the relatively poorer DoTCP performance of Cloudflare, Google, and OpenDNS resolvers for requests originating from Africa over IPv6. Analyzing DNS response times over IPv6 by Autonomous System reveals higher variations compared to IPv4. Generally, probes from UUNET (21.1ms) and KPN (15.6ms) receive the fastest DoTCP responses, with most Public resolvers displaying their minimal DoTCP response times for these two ASes. CleanBrowsing shows outliers in DoTCP response times for probes from the Orange (191.4ms), UUNET (296.6ms), and KPN (349.5ms). UncensoredDNS shows relatively high response times for requests from all Autonomous Systems (62.2ms-214.1ms) except for UUNET

(20.5ms).

Takeaway: In IPv4, Cloudflare and Google DNS resolvers demonstrate the most stable DoTCP and DoUDP response times across all continents. Other Public resolvers generally have significantly higher DoTCP response times for at least one, and often multiple, continents, particularly in Africa (322.5ms) and South America (245.8ms) where they take the longest to respond. Similar patterns are observed for DoUDP response times. Whereas, in IPv6 measurements, Europe exhibits the lowest median response times for DoTCP and DoUDP, while Africa exhibits the highest. Cloudflare seems to have fewer IPv6-capable Points-of-Presence distributed in Africa, resulting in increased round-trip times (RTTs) due to long distances.

Combination - IPv4 and IPv6. The evaluation of the measured response times of DNS over both transport protocols and IP versions as CDF emphasizes the results summarized above and allows their comparison from a different standpoint. Instead of the median aggregated median RTs of each probe shown in Fig. [5-6], Fig. 7 shows the accumulation of all recorded probe-medians based on the transport protocol and IP version used. the Fig. 7a confirms the extremely high DoTCP response times of Neustar DNS over IPv4. It shows that around 80% of the DoTCP requests take more than one second to be answered. The comparably bad performance of Yandex and UncensoredDNS can be seen for all combinations of transport protocol and IP version, especially for UncensoredDNS over UDP and IPv4 (more than 90% of the requests have an RT of more than 150ms). Furthermore, the response times of CleanBrowsing over both transport protocols on average increase when IPv6 is used instead of IPv4. This observation is not reflected by the median response times as the effect shows mainly for the slowest 25% of requests. Figure 7 also confirms that Cloudflare and Google exhibit the most stable DNS response times of all Public resolvers.

3) **Traceroute:** RIPE Atlas does not allow Traceroute measurements within the probe networks, and the focus of this paper is primarily on the performance, resilience, and security of Public resolvers while Probe resolvers are excluded from the subsequent analysis. Ideally, the resolvers should exhibit DoUDP response times that are roughly similar to the RTTs measured through Traceroute; however, not exactly the same, as ICMP (used in Traceroute) requires slightly less overhead for transportation than UDP. Fig. 8 presents the cumulative distribution function (CDF) of median response times and round-trip times for all Public resolvers over each probe, validating these expectations for both IPv4 and IPv6. DoTCP over IPv6 aligns more closely with the anticipated ratio of response times to round-trip times. The response time/round-trip time ratio (RT/RTT ratio) represents the quotient of the median response time and the median RTT between the probe and the resolver. Ideally, resolvers should have an RT/RTT ratio of 1 for DoUDP and 2 for DoTCP. Fig. 9 presents the RT/RTT ratios per resolver, transport protocol, and IP version. For Neustar over DoTCP and IPv4, the observed ratio exceeds 2 for over 80% of the probes. This further highlights that

TABLE IV: EDNS(0) Buffer Sizes announced to RIPE probes by public resolvers in the measurement series from the core. Buffer sizes that are not equal to 512, 1232, or 4096 bytes are summarized in the column other. If EDNS is not used at all this is reflected in the column none.

		512	1232	4096	none	other
CleanBrowsing	IPv4	97.83%	1.04%	0.56%	0.31%	0.27%
	IPv6	99.33%	0.19%	0.20%	0.29%	0.00%
Cloudflare	IPv4	0.47%	98.24%	0.61%	0.31%	0.38%
	IPv6	0.22%	99.29%	0.20%	0.28%	0.00%
Comodo	IPv4	0.45%	1.02%	97.94%	0.31%	0.27%
	IPv6	-	-	-	-	-
Google	IPv4	97.61%	1.20%	0.53%	0.30%	0.35%
	IPv6	99.23%	0.18%	0.31%	0.29%	0.00%
Neustar	IPv4	0.45%	0.98%	97.99%	0.31%	0.27%
	IPv6	0.24%	0.18%	99.29%	0.29%	0.00%
OpenDNS	IPv4	0.54%	0.98%	97.86%	0.31%	0.31%
	IPv6	0.24%	0.18%	99.30%	0.29%	0.00%
OpenNIC	IPv4	0.46%	98.41%	0.56%	0.31%	0.27%
	IPv6	0.24%	99.27%	0.20%	0.29%	0.00%
Quad9	IPv4	27.12%	71.66%	0.57%	0.32%	0.33%
	IPv6	25.78%	73.73%	0.20%	0.29%	0.00%
UncensoredDNS	IPv4	1.84%	93.55%	2.27%	1.25%	1.09%
	IPv6	0.37%	98.95%	0.28%	0.40%	0.00%
Yandex	IPv4	0.55%	1.02%	97.68%	0.47%	0.27%
	IPv6	0.23%	0.18%	99.30%	0.29%	0.00%
Overall	IPv4	27.41%	35.66%	36.24%	0.36%	0.33%
	IPv6	26.04%	39.39%	34.27%	0.30%	0.00%

the very high response times of the resolver are not due to a sparsely distributed global network of Points-of-Presence but rather an inadequate DoTCP implementation at various PoPs. As explained in existing literature [46][47], an inadequate DoTCP implementation generally refers to an inappropriate or insufficient configuration of the DoTCP protocol along with a failure to update the DNS-related software. This could happen as a result of not adhering to the standard configuration while implementing DoTCP.

The same conclusion can be drawn for UncensoredDNS over both TCP and UDP, as shown in Fig. 9. OpenDNS and Google exhibit higher RT/RTT ratios compared to other Public resolvers over both transport protocols and IP versions, indicating that their relatively fast DNS response times are more attributed to a well-distributed global network rather than exceptionally efficient request processing.

C. Evaluation from the core

1) **Failure rate:** Fig. 10 exhibits that the failure rates for DoTCP requests over IPv4 from the core using public resolvers are higher in the current measurement series (9.09%) compared to previous measurements (4.01%). CleanBrowsing, Cloudflare, Google, OpenDNS, OpenNIC, and Yandex demonstrate high resiliency with DoTCP failure rates ranging from 1.34% to 2.62%. Similarly, Quad9 shows higher failure rates from multiple regions and ASs. When using IPv6, the overall failure rates for all public resolvers slightly increase to 11.53%. More READ-errors (19.67%) and fewer Timeouts (29.99%) are observed compared to measurements from the edge. Quad9 and UncensoredDNS exhibit the highest failure percentages. Fig. 11 shows that from all continents, the DoTCP failure rates for the public resolvers is surged in comparison to that of explained in Fig. 3. However, more than 88% of the

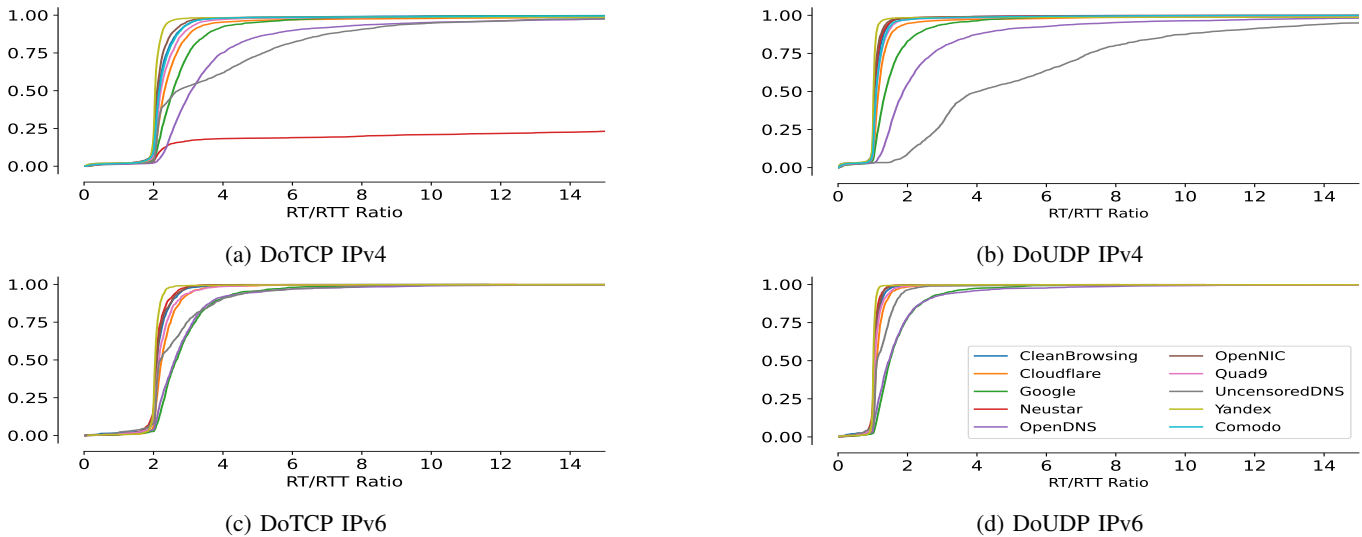


Fig. 9: RT-RTT ratio (quotient of median round-trip-time and median response time for each probe) for each resolver over TCP/UDP and IPv4/IPv6 as CDF.

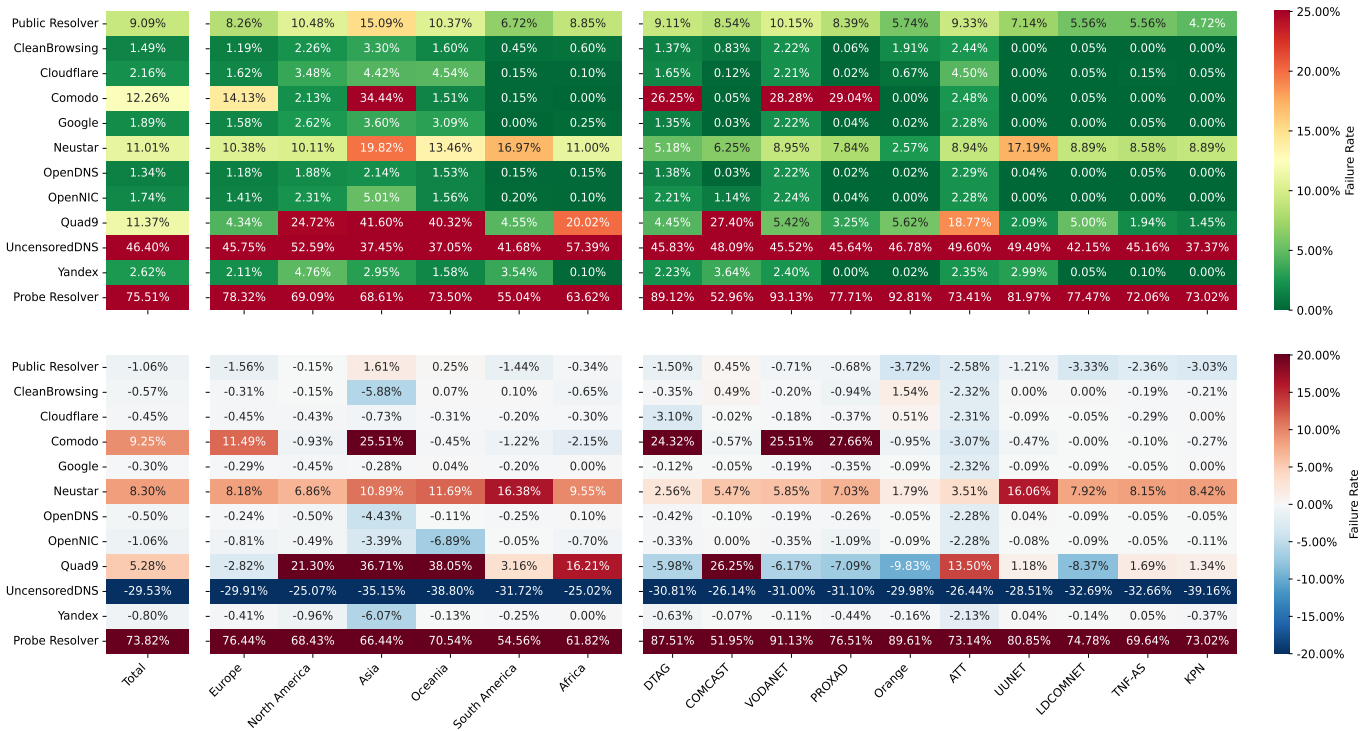


Fig. 10: Failure rates observed from the core over IPv4. The upper part represents the DoTCP failure rates of all resolvers in total and per continent and AS. The lower part reflects the difference between the DoTCP and the DoUDP failure rates for a particular pairing (a negative value hints at a higher DoUDP failure rate). Public Resolver summarizes the observations of all resolvers that are not Probe resolvers.

measurements to the public resolvers over both IP versions and transport protocols receive a valid DNS response. This shows that there are no significant problems with our NSEs and yields enough data to reliably analyze the resolvers' DoTCP usage and EDNS(0).

2) **EDNS(0) Configurations:** The buffer sizes that the resolvers advertise in the case of cached domains are expected

to be similar to the ones observed in the measurements before as the only change that has been made on the client side is that TXT records are requested instead of A records. Table IV confirms these expectations showing small differences between IPv4 and IPv6 for all resolvers and similar usage rates. Again, CleanBrowsing and Google mostly advertise a UDP buffer size of 512 bytes. To examine the usage of EDNS(0) configurations

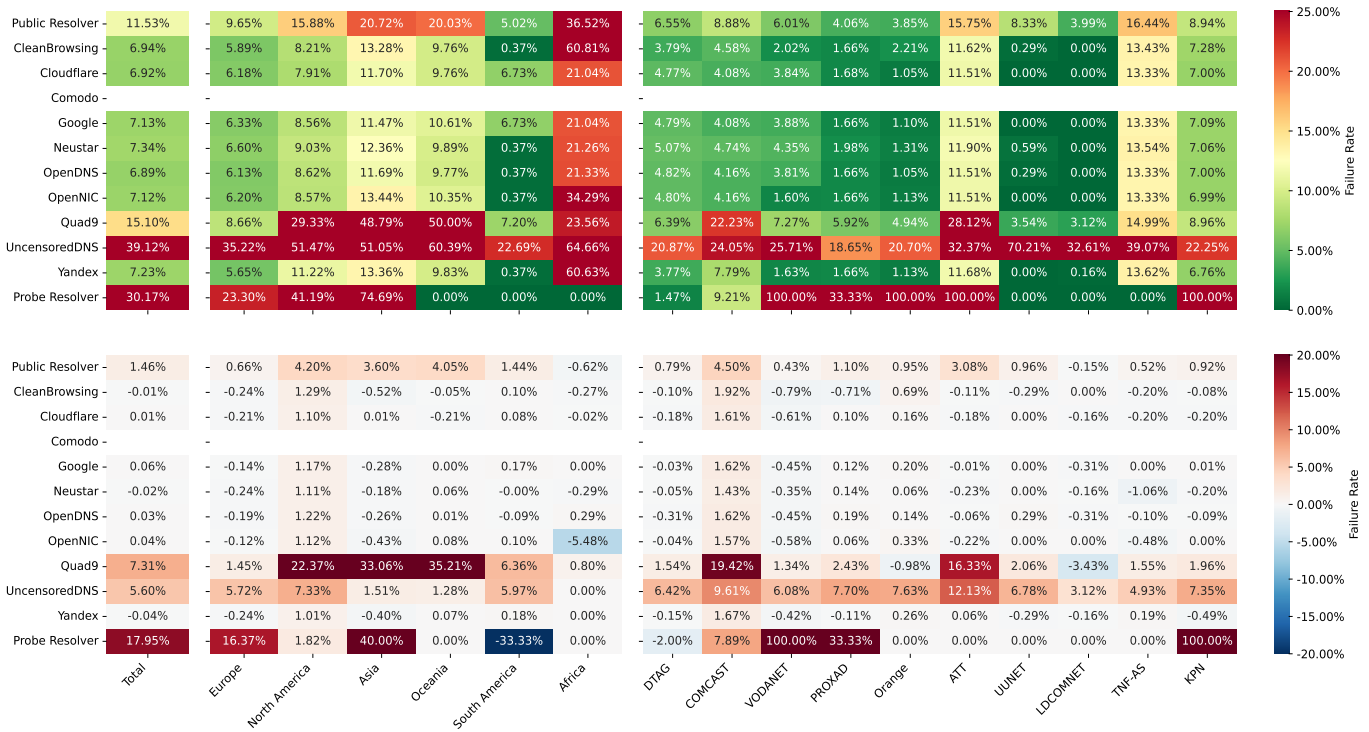


Fig. 11: Failure rates observed from the core over IPv6. The upper part presents the failure rates over DoTCP, and the lower one is the difference between DoTCP and DoUDP failure rates. White cells indicate that there is no data for the given pairing.

TABLE V: Distribution of the resolvers communicating with the authoritative name servers for uncached domains used by the public resolvers over AS.

		Choopa	Cogent	Packet	Cloudfl.	Google	Neust.	O-DNS	Myth.	W-net-1	FSKNET	Yandex	Other
CleanBrowsing	IPv4	66.58%	28.54%	-	0.64%	0.45%	-	0.05%	-	0.05%	-	-	3.69%
	IPv6	59.07%	14.89%	25.26%	0.07%	0.21%	-	-	-	-	-	-	0.49%
Cloudflare	IPv4	0.06%	0.04%	-	98.04%	0.46%	-	0.04%	-	0.05%	-	-	1.32%
	IPv6	-	-	-	99.38%	0.21%	-	-	-	-	-	-	0.41%
Comodo	IPv4	95.63%	0.04%	-	0.67%	0.48%	-	0.05%	-	0.06%	-	-	3.07%
	IPv6	-	-	-	-	-	-	-	-	-	-	-	-
Google	IPv4	0.06%	0.04%	-	0.64%	97.86%	-	0.05%	-	0.09%	-	-	1.28%
	IPv6	-	-	-	0.06%	99.41%	-	-	-	-	-	-	0.52%
Neustar	IPv4	0.06%	0.04%	-	0.63%	0.48%	94.79%	0.05%	-	0.06%	-	-	3.90%
	IPv6	-	-	-	0.06%	0.23%	99.23%	-	-	-	-	-	0.48%
OpenDNS	IPv4	0.06%	0.04%	-	0.59%	0.53%	-	97.68%	-	0.01%	-	-	1.10%
	IPv6	-	-	-	0.07%	0.22%	-	99.30%	-	-	-	-	0.41%
OpenNIC	IPv4	87.92%	0.03%	-	0.59%	0.45%	-	0.05%	9.77%	0.05%	-	-	1.12%
	IPv6	65.21%	-	-	0.06%	0.23%	-	-	34.09%	-	-	-	0.41%
Quad9	IPv4	0.06%	0.04%	-	0.70%	0.51%	-	0.05%	-	82.34%	-	-	16.30%
	IPv6	-	-	-	0.08%	0.22%	-	-	-	75.57%	-	-	24.13%
UncensoredDNS	IPv4	0.15%	0.09%	-	1.49%	1.14%	-	0.12%	-	0.14%	94.07%	-	2.81%
	IPv6	-	-	-	0.08%	0.34%	-	-	-	-	98.98%	-	0.60%
Yandex	IPv4	0.06%	0.04%	-	0.69%	0.56%	-	0.04%	-	0.05%	-	97.43%	1.13%
	IPv6	-	-	-	0.06%	0.22%	-	-	-	-	-	99.30%	0.41%
Overall	IPv4	18.87%	3.52%	-	12.35%	12.22%	10.98%	11.83%	1.19%	9.19%	4.54%	11.75%	3.57%
	IPv6	14.49%	1.72%	2.95%	11.50%	11.70%	11.44%	11.51%	3.97%	8.33%	7.85%	11.46%	3.09%

by resolvers communicating with authoritative NSes for uncached domains, we present a distribution overview in Table V. Most resolvers exhibit a preferred AS, with Cloudflare, Google, Neustar, OpenDNS, and Yandex DNS primarily using their own ASs for over 94% of resolutions. Public resolvers universally employ DoUDP, emphasizing the importance of proper EDNS(0) buffer sizes in the core. Commonly used buffer sizes include 1400, 1410, and 1452 bytes. Additional buffer size details are displayed in Table VI.

3) **EDNS Options:** In Table VII, a comprehensive listing of options employed by DNS resolvers in communication with name servers is provided. EDNS(0) is utilized by all DNS

resolvers in the majority of cases (>99.84%). Among the advertised options by public resolvers, Cookie (4.80% IPv4, 7.91% IPv6) and EDNS Client Subnet (ECS) (1.81% IPv4, 1.49% IPv6) are notable but Google predominantly employs ECS (14.24% IPv4, 12.53% IPv6). However, other options including client subnet information, are transmitted in less than 0.24% of requests. Callejo *et al.* in [13] show only ECS caching support by Google >5%. We extend our studies to explore other EDNS options, such as cookie options, as well. RFC 7871 [48] specifies that NSes should include ECS with matching parameters in their response. Google indicates that if name servers do not support ECS, Google public DNS may

TABLE VI: EDNS(0) buffer sizes announced to the authoritative NSes. Other buffer sizes and cases in EDNS that are not used are summarized in the column "other". NOTE: CB= CleanBrowsing; U-DNS= UncensoredDNS. All the values are in percentage (%).

		512.0	1232.0	1400.0	1410.0	1452.0	4096.0	other
CB	IPv4	0.11	98.24	0.45	0.05	0.64	0.36	0.16
	IPv6	0.01	99.47	0.21	0.00	0.07	0.05	0.19
Cloudflare	IPv4	0.36	0.65	0.46	0.04	98.04	0.30	0.16
	IPv6	0.01	0.26	0.21	0.00	99.38	0.04	0.10
Comodo	IPv4	0.11	0.70	0.48	0.05	0.67	95.21	2.78
	IPv6	-	-	-	-	-	-	-
Google	IPv4	0.22	0.78	97.86	0.05	0.64	0.27	0.19
	IPv6	0.02	0.26	99.41	0.00	0.06	0.14	0.10
Neustar	IPv4	0.04	0.70	0.48	0.05	0.63	97.45	0.66
	IPv6	0.02	0.31	0.23	0.00	0.06	98.79	0.60
OpenDNS	IPv4	0.08	0.61	0.53	97.68	0.59	0.32	0.19
	IPv6	0.01	0.26	0.22	99.30	0.07	0.04	0.10
OpenNIC	IPv4	0.06	98.29	0.45	0.05	0.59	0.37	0.18
	IPv6	0.01	99.56	0.23	0.00	0.06	0.05	0.10
Quad9	IPv4	0.07	98.05	0.51	0.05	0.70	0.40	0.21
	IPv6	0.01	99.54	0.22	0.00	0.08	0.04	0.11
U-DNS	IPv4	2.68	93.15	1.14	0.12	1.49	0.97	0.45
	IPv6	1.46	97.91	0.34	0.00	0.08	0.07	0.15
Yandex	IPv4	0.03	0.65	0.56	0.04	0.69	92.86	5.16
	IPv6	0.00	0.26	0.22	0.00	0.06	94.31	5.14
Overall	IPv4	0.24	39.74	12.22	11.83	12.34	22.78	0.85
	IPv6	0.13	42.09	11.70	11.51	11.49	22.33	0.75

TABLE VII: EDNS options announced to the authoritative name servers

		EDNS	Cookie	ECS
CleanBrowsing	IPv4	99.93%	0.22%	0.10%
	IPv6	99.91%	0.05%	0.04%
Cloudflare	IPv4	99.94%	0.32%	0.10%
	IPv6	100.00%	0.05%	0.05%
Comodo	IPv4	98.10%	0.33%	0.11%
	IPv6	-	-	-
Google	IPv4	99.93%	0.31%	14.23%
	IPv6	100.00%	0.16%	12.53%
Neustar	IPv4	99.93%	0.23%	0.10%
	IPv6	99.93%	0.05%	0.04%
OpenDNS	IPv4	99.94%	0.22%	0.10%
	IPv6	100.00%	0.05%	0.04%
OpenNIC	IPv4	99.93%	0.22%	0.11%
	IPv6	100.00%	0.05%	0.05%
Quad9	IPv4	99.93%	0.24%	0.13%
	IPv6	100.00%	0.06%	0.03%
UncensoredDNS	IPv4	99.84%	94.62%	0.24%
	IPv6	100.00%	99.06%	0.06%
Yandex	IPv4	99.93%	0.22%	0.11%
	IPv6	100.00%	0.05%	0.04%
Overall	IPv4	99.93%	4.80%	1.81%
	IPv6	99.98%	7.91%	1.49%

refrain from sending ECS queries to them. This suggests that Google's usage of ECS could be higher if servers appropriately handle the requests. Nevertheless, multiple Google resolvers within the core, identifiable by their IP addresses, transmit subnet information to our servers. The question of Google's ECS usage rate when communicating with NSes that correctly respond remains open for further investigation.

TABLE VIII: Percentage of measurements that were denoted as successful, but did not receive an answer section.

	Answer Received
CleanBrowsing	1.55%
Cloudflare	1.63%
Comodo	62.79%
Google	57.17%
Neustar	61.37%
OpenDNS	97.65%
OpenNIC	43.39%
Quad9	43.02%
UncensoredDNS	4.05%
Yandex	64.44%
Total	46.48%

TABLE IX: Failure Rates of the RIPE Atlas measurements from the core when the unique domain generated for the request was never requested. NR = Never Requested

		Failure	NR, Failure	NR, Success
CleanBrowsing	IPv4	1.07%	2.13%	0.42%
	IPv6	1.56%	6.53%	0.56%
Cloudflare	IPv4	0.56%	2.16%	0.86%
	IPv6	1.42%	6.20%	0.81%
Comodo	IPv4	24.25%	5.11%	1.27%
	IPv6	-	-	-
Google	IPv4	2.92%	1.95%	0.91%
	IPv6	29.74%	6.74%	0.83%
Neustar	IPv4	6.19%	2.69%	0.31%
	IPv6	33.01%	6.83%	0.18%
OpenDNS	IPv4	3.88%	1.83%	0.62%
	IPv6	36.61%	6.64%	0.46%
OpenNIC	IPv4	0.47%	97.31%	0.16%
	IPv6	0.06%	99.31%	0.01%
Quad9	IPv4	7.53%	8.91%	0.69%
	IPv6	20.71%	6.69%	0.69%
UncensoredDNS	IPv4	0.60%	85.85%	0.19%
	IPv6	1.18%	10.07%	1.07%
Yandex	IPv4	4.47%	3.78%	0.16%
	IPv6	5.14%	6.75%	0.11%
All	IPv4	5.19%	21.16%	0.56%
	IPv6	14.40%	17.27%	0.52%

Takeaway: Almost 40% of public resolvers announce very large EDNS(0) buffer sizes (4096B) both from the Edge as well as from the Core, which potentially causes fragmentation.

4) **Valid/ Invalid responses:** As it was observed that the transport protocol used by the probes does not affect the usage of DoTCP/ DoUDP in the communication between resolvers and the NSes, all measurements in this experiment are carried out over DoUDP. Overall, 11,637,539 individual measurements are conducted based on unique domain names. We furthermore observe that the NS returning 2KB responses receives more requests (5,642,439) than the one returning 4KB (2,395,455). Moreover, some domains are requested on both NSes (2,733,540 results). It is important to note that in many cases where a successful DNS response was indicated, the response did not contain an answer section (in 53.52% of cases) as presented in Table VIII. Among these cases, the vast majority (97.35%) were truncated responses (TC-bit set), while 2.22% denoted a server failure. For the remaining

0.43%, no clear reason for the missing answer section could be identified. The truncation of responses can be attributed to the limitation imposed by RIPE Atlas on UDP buffer sizes, which is set to a maximum of 4096 bytes which explains the high amount of truncated responses in general. OpenDNS consistently provided a valid answer section in most cases (97.65%), while others included an answer section in less than 64.44% of their responses. Surprisingly, RIPE Atlas measurements ended successfully even after receiving a response with the TC-bit set, indicating a lack of proper fallback to DoTCP in many probes. In Table IX these cases are taken into consideration and the respective failure rates of the RIPE Atlas measurements are presented. Additionally, there were cases where certain domains were never requested at any of the servers, contributing to a failed DNS response (21.16% IPv4, 17.27% IPv6). Other resolvers had a failure rate of over 6.53% for IPv6 requests that were not forwarded to authoritative name servers. This behavior may be attributed to some resolvers blacklisting our authoritative name server due to the receipt of large responses.

5) **Canonical/ Non-canonical requests:** We begin our analysis by classifying incoming requests as canonical and non-canonical according to Mao *et al.*'s work [49]. We then evaluate the DoTCP usage rates of the resolvers to assess their response to large buffer sizes. Additionally, we introduce a scenario involving a single incoming UDP request. To focus on the resolvers' reactions to response sizes of 2KB and 4KB, we only consider results that can be directly matched to one of the servers and their respective response sizes. Table X displays the resolvers' usage of different scenarios when communicating with the 2KB name server. Notably, CleanBrowsing, Cloudflare, Google, OpenDNS, and UncensoredDNS predominantly send a UDP message followed by a TCP message. As indicated in Table VI, the resolvers advertise EDNS(0) buffer sizes of 1452B or less, demonstrating the expected fallback behavior. Table X presents the usage of different scenarios by the resolvers in response to 4KB name server replies. Quad9 demonstrates more non-canonical responses to 4KB compared to 2KB responses.

6) **TCP Usage:** To assess TCP usage, we examined the presence of DoTCP requests within the query sequence reaching the name servers. Table XI shows the DoTCP usage rates of resolvers when receiving 2KB responses. Resolvers primarily employing canonical scenarios consistently utilize TCP in their final request, including Quad9 (99.69% IPv4, 99.70% IPv6). Yandex and Comodo rarely use DoTCP with 2KB responses in the last request. When receiving 4KB responses (see Table XI), almost all resolvers employ TCP in the majority of measurements for both IP versions (>98.67%). However, a notable number of measurements lack TCP usage by several resolvers (up to 1.33%), indicating possible fragmentation between the name server and resolver.

Takeaway: Each resolver exhibits one preferred buffer size which is advertised to the name servers in more than 90% of the cases.

TABLE X: Classification of incoming sequences of DNS queries at the 2KB and 4KB name server for each resolver.

Resolvers		2KB			4KB		
		Canonical	Non-Canonical	Single UDP	Canonical	Non-Canonical	Single UDP
CleanBrowsing	IPv4	99.42%	0.47%	0.11%	98.65%	1.27%	0.08%
	IPv6	99.56%	0.24%	0.21%	98.83%	0.25%	0.92%
Cloudflare	IPv4	95.36%	4.45%	0.19%	91.76%	8.22%	0.00%
	IPv6	91.96%	7.70%	0.34%	85.24%	14.37%	0.36%
Comodo	IPv4	1.19%	97.83%	0.98%	8.91%	91.07%	0.02%
	IPv6	-	-	-	-	-	-
Google	IPv4	99.24%	0.67%	0.10%	98.67%	0.83%	0.51%
	IPv6	99.44%	0.26%	0.30%	98.82%	0.19%	0.99%
Neustar	IPv4	1.31%	97.61%	1.08%	3.13%	96.86%	0.01%
	IPv6	0.25%	98.67%	1.07%	0.73%	98.25%	1.02%
OpenDNS	IPv4	97.71%	2.08%	0.21%	97.55%	2.37%	0.08%
	IPv6	99.13%	0.62%	0.25%	98.38%	0.79%	0.83%
OpenNIC	IPv4	55.68%	35.60%	8.72%	64.36%	35.42%	0.22%
	IPv6	35.20%	26.89%	37.91%	38.34%	7.06%	54.60%
Quad9	IPv4	46.14%	53.67%	0.19%	26.37%	73.60%	0.03%
	IPv6	46.28%	53.47%	0.25%	31.23%	68.21%	0.56%
UncensoredDNS	IPv4	94.11%	4.70%	1.19%	94.52%	5.43%	0.05%
	IPv6	99.52%	0.20%	0.29%	99.48%	0.11%	0.41%
Yandex	IPv4	1.19%	97.79%	1.02%	3.34%	96.56%	0.10%
	IPv6	0.22%	98.94%	0.84%	0.83%	97.86%	1.31%
All	IPv4	59.69%	39.82%	0.48%	70.91%	28.95%	0.13%
	IPv6	69.05%	30.50%	0.45%	75.43%	23.79%	0.78%

TABLE XI: TCP usage of DNS resolvers when 2KB and 4KB responses are received. TCP Used represents all scenarios in which TCP is used at any point in the request sequence. For Last TCP, only those sequences ending with a DoTCP request are considered.

Resolvers		2KB		4KB	
		TCP Used	Last TCP	TCP Used	Last TCP
CleanBrowsing	IPv4	99.84%	99.80%	99.92%	99.77%
	IPv6	99.76%	99.76%	99.08%	99.02%
Cloudflare	IPv4	99.74%	96.95%	100.00%	95.76%
	IPv6	99.60%	95.84%	99.64%	92.63%
Comodo	IPv4	7.94%	3.36%	99.98%	99.52%
	IPv6	-	-	-	-
Google	IPv4	99.86%	99.81%	99.49%	99.42%
	IPv6	99.65%	99.65%	99.00%	98.97%
Neustar	IPv4	73.52%	49.96%	99.99%	99.78%
	IPv6	72.17%	48.46%	98.98%	98.91%
OpenDNS	IPv4	99.73%	99.67%	99.91%	99.81%
	IPv6	99.70%	99.70%	99.16%	99.13%
OpenNIC	IPv4	88.05%	85.16%	99.78%	94.72%
	IPv6	54.35%	54.35%	45.09%	41.72%
Quad9	IPv4	99.74%	99.69%	99.97%	99.80%
	IPv6	99.70%	99.70%	99.43%	99.37%
UncensoredDNS	IPv4	98.34%	98.04%	99.95%	99.21%
	IPv6	99.66%	99.66%	99.59%	99.58%
Yandex	IPv4	4.49%	3.17%	99.85%	99.54%
	IPv6	1.58%	0.94%	98.67%	98.58%
All	IPv4	75.36%	71.97%	99.86%	98.79%
	IPv6	84.25%	81.38%	99.22%	97.84%

V. IMPLICATIONS ON NETWORK AND SERVICE MANAGEMENT

The discussion in this paper revolves around the profound implications of network and service management to ideally empower the network service providers in making informed decisions to better manage their critical DNS services.

DNS, traditionally overseen by ISPs, has witnessed an emerging management trend with cloud providers like Google and Cloudflare actively participating. Our paper presents results, particularly influenced by the DNS Flag Day 2020, offering insights to these stakeholders for efficiently managing their crucial Internet services. The study has focused on evaluating the resiliency of DNS infrastructure amidst the escalating DNS response sizes causing truncation and fragmentation. As a result, we gain key insights into ensuring robust and reliable DNS network services. Notably, this study also investigated

the capability of DNS to fallback to a more reliable transport protocol (DNS-over-TCP) in the face of fragmentation and truncation. We conduct a comprehensive performance evaluation, assessing the responsiveness of DoTCP in comparison to traditional DNS-over-UDP. The results outlined in this study would substantially facilitate the standards bodies, like the IETF, to evaluate the implications of transitioning to a more reliable transport protocol, preferably, DNS-over-TCP as of now and DNS-over-QUIC in the near future. Towards the end, we also discuss a few emerging technologies and standards in depth (with EDNS, and ECS options).

Takeaway: *In most case, while responding to 2KB and 4KB responses from ANSes, resolvers like Yandex, Neustar, Comodo, and Quad9, do not fall back to DoTCP. This bears the risk of fragmented responses.*

VI. LIMITATIONS AND FUTURE WORK

Approximately 88% of our probe measurements are concentrated in North America and Europe, limiting the generality of DNS resiliency observations to other regions. To address this limitation, we provide response times categorized by continent. However, it is important to note that observations for continents with fewer probes have smaller sample sizes, which hinders drawing reliable conclusions. Similarly, when analyzing response times of specific autonomous systems, particularly over IPv6, the sample size remains relatively low. The study of EDNS(0) options focuses on the communication between different resolvers and our custom authoritative NSes. Therefore, the usage numbers may not accurately represent the capabilities of the resolvers and their EDNS(0) options in general. The observations reveal various non-canonical sequences employed by DNS resolvers in response to large response sizes. Further investigation is required to fully understand the behavior of different resolvers, including their adjustment of announced EDNS(0) buffer sizes when receiving large responses.

While our study emphasized the unencrypted DNS protocols DoUDP and DoTCP, the recently standardized encrypted DNS protocol DNS-over-QUIC (DoQ) (RFC 9250)[50][51][52][53] does inherently solve fragmentation through the QUIC protocol (RFC 9000)[54] while also supporting increased DNS message sizes. Though some studies show that adoption of DoQ is scarce [55][56], there is a separate branch of literature [51][52] highlighting the benefits of implementing DoQ in different workloads compared to the other encrypted DNS protocols. This reflects the relevance of DoQ in the current world. We, therefore, believe that DoQ is supposedly going to become the predominant successor amongst the other encrypted DNS protocols, thereby warranting a detailed investigation of DoQ in the future.

VII. DISCUSSION

A. Geolocation of the Resolvers

The primary objective of this study has been to demonstrate the responsiveness of DoTCP and DoUDP over IPv4 and IPv6. Therefore, performing a geolocation validation for anycast and

unicast resolvers is something we did not consider in the current work, as it seemed orthogonal to our objective. Moreover, certain studies [57–59] in the past have demonstrated that geolocation services often fail to provide precise location information. We also realize that a sparsely distributed global network of Points-of-Presence is the major reason behind the high response time of certain resolvers. Though limiting the number of probes with equal geographical distribution might seem to be a viable solution in this case, it can lead to a reduction in data volume. This will vastly impact the analysis of our results while generalizing the performance of the protocols (DoTCP, DoUDP). It would be more appropriate to employ a limited number of probes with equal geographical distribution to prevent bias when looking at a narrow region, such as country-level or state-level performance analysis.

B. Routing and Stacking

The handling of TCP and UDP traffic by routers, including load balancing decisions like ECMP [60], can affect response times. Other factors like network congestion, routing policies, and quality-of-service (QoS) [61] may also contribute to variations. Further investigation into router load-balancing strategies for different protocols is necessary for a deeper understanding of resolver performance. The TCP stack [62] can also impact connection speed, although studying its effects would require significant modifications to the experimental setup. Exploring factors like enabled SYN cookies in the TCP stack [6] and their impact on connection establishment time could be considered in future studies.

VIII. CONCLUSION

We conducted the measurements analyzing DoTCP resiliency, responsiveness and deployment from the edge and the core over IPv4 and IPv6. Additionally, the EDNS(0) configurations of ten public resolvers were studied. Issuing more than 14M individual DNS requests using 2527 globally distributed RIPE Atlas probes, we performed multiple experiments focusing on observations to conclude that most resolvers show similar resiliency for both DoTCP and DoUDP where 3 out of 10 resolvers mainly announce very large EDNS(0) buffer sizes, which potentially causes fragmentation. The analysis of DoTCP and DoUDP performance revealed significant regional variations for both IP versions. Notably, requests originating from Africa or South America exhibited the highest median response times. This highlights the need for further investigation and optimization in such regions. Particularly over IPv4, Cloudflare and Google emerged as the Public resolvers with the most consistent and stable response times across all continents. In reaction to large response sizes from authoritative name servers, we find that resolvers do not fall back to the usage of DoTCP in many cases, bearing the risk of fragmented responses. As the message sizes in the DNS are expected to grow further, this problem will become more urgent in the future.

ACKNOWLEDGMENT

This work was supported by the Volkswagenstiftung Niedersächsisches Vorab (Funding No. ZN3695).

REFERENCES

- [1] Paul V. Mockapetris and Kevin J. Dunlap. Development of the domain name system. In Vinton G. Cerf, editor, *Proceedings of the ACM Symposium on Communications Architectures and Protocols, SIGCOMM 1988, Stanford, CA, USA, August 16-18, 1988*, pages 123–133. ACM, 1988.
- [2] Paul V. Mockapetris. Domain names - concepts and facilities. *RFC*, 1034:1–55, 1987.
- [3] Jon Postel. Transmission control protocol. *RFC*, 793:1–91, 1981.
- [4] Jon Postel. User datagram protocol. *RFC*, 768:1–3, 1980.
- [5] Paul V. Mockapetris. Domain names - implementation and specification. *RFC*, 1035:1–55, 1987.
- [6] John Dickinson, Sara Dickinson, Ray Bellis, Allison Mankin, and Duane Wessels. DNS transport over TCP - implementation requirements. *RFC*, 7766:1–19, 2016.
- [7] Ray Bellis. DNS transport over TCP - implementation requirements. *RFC*, 5966:1–7, 2010.
- [8] Paul Vixie. Extension mechanisms for DNS (EDNS0). *RFC*, 2671:1–7, 1999.
- [9] Ivica Stipovic. Analysis of an extension dynamic name service - A discussion on DNS compliance with RFC 6891. *CoRR*, abs/2003.13319, 2020.
- [10] Axel Koolhaas and Tjeerd Slokker. Defragmenting DNS: Determining the Optimal Maximum UDP Response Size for DNS, 2020. [Last Accessed: 19.April.2023]: <https://bit.ly/3Ag6Mck>.
- [11] Mike Kosek, Trinh Viet Doan, Simon Huber, and Vaibhav Bajpai. Measuring DNS over TCP in the Era of increasing DNS Response Sizes: A View from the Edge. *Computer Communication Review*, 2022.
- [12] Vaibhav Bajpai and Jürgen Schönwälder. A Longitudinal View of Dual-Stacked Websites - Failures, Latency and Happy Eyeballs. *IEEE/ACM Trans. Netw.*, 27(2):577–590, 2019.
- [13] Patricia Callejo, Marcelo Bagnulo, Jaime González Ruiz, Andra Lutu, Alberto García-Martínez, and Rubén Cuevas. Measuring doh with web ads. *Comput. Networks*, 212:109046, 2022.
- [14] Pratyush Dikshit, Mike Kosek, Nils Faulhaber, Jayasree Sengupta, and Vaibhav Bajpai. Evaluating DNS Resiliency with Truncation, Fragmentation and DoTCP Fallback. In *2023 IFIP Networking Conference (IFIP Networking)*, 2023.
- [15] Giovane C. M. Moura, Moritz Müller, Marco Davids, Maarten Wullink, and Cristian Hesselman. Fragmentation, Truncation, and Timeouts: Are Large DNS Messages Falling to Bits? In *Passive and Active Measurement Conference (PAM)*, 2021.
- [16] Ron Bonica, Fred Baker, Geoff Huston, Robert M. Hinden, Ole Troan, and Fernando Gont. IP fragmentation considered fragile. *RFC*, 8900:1–23, 2020.
- [17] Amir Herzberg and Haya Shulman. Fragmentation considered poisonous, or: One-domain-to-rule-them-all.org. In *IEEE Conference on Communications and Network Security (CNS) 2013*. IEEE, 2013.
- [18] Haya Shulman and Michael Waidner. Fragmentation Considered Leaking: Port Inference for DNS Poisoning. In *Applied Cryptography and Network Security Conference, (ACNS)*. Springer, 2014.
- [19] Soolee Son and Vitaly Shmatikov. The hitchhiker's guide to DNS cache poisoning. In Sushil Jajodia and Jianying Zhou, editors, *Security and Privacy in Communication Networks - 6th International ICST Conference, SecureComm 2010, Singapore, September 7-9, 2010. Proceedings*, volume 50 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 466–483. Springer, 2010.
- [20] Harel Berger, Amit Z. Dvir, and Moti Geva. A wrinkle in time: A case study in DNS poisoning. *International Journal of Information Security*, 20(3):313–329, 2021.
- [21] Chad R Dougherty. Multiple dns implementations vulnerable to cache poisoning, <https://www.kb.cert.org/vuls/id/800113>, 2008.
- [22] Matt Crawford. Transmission of IPv6 Packets over Ethernet Networks. *RFC*, 2464:1–7, 1998.
- [23] Jin Cao, Maode Ma, Xilei Wang, and Haochen Liu. A selective re-query case sensitive encoding scheme against DNS cache poisoning attacks. *Wireless Personal Communications*, 94(3):1263–1279, 2017.
- [24] Tianxiang Dai, Haya Shulman, and Michael Waidner. Dns-over-tcp considered vulnerable. In *ANRW '21: Applied Networking Research Workshop, Virtual Event, USA, July 24-30, 2021*, pages 76–81. ACM, 2021.
- [25] Nicholas Weaver, Christian Kreibich, Boris Nechaev, and Vern Paxson. Implications of netalyzr's dns measurements. In *Proceedings of the First Workshop on Securing and Trusting Internet Names (SATIN), Teddington, United Kingdom*, 2011.
- [26] Zi Hu, Liang Zhu, John S. Heidemann, Allison Mankin, Duane Wessels, and Paul E. Hoffman. Specification for DNS over transport layer security (TLS). *RFC*, 7858:1–19, 2016.
- [27] Paul E. Hoffman and Patrick McManus. DNS queries over HTTPS (doh). *RFC*, 8484:1–21, 2018.
- [28] Pratyush Dikshit, Jayasree Sengupta, and Vaibhav Bajpai. Recent trends on privacy-preserving technologies under standardization at the IETF. *Comput. Commun. Rev.*, 53(2):22–30, 2023.
- [29] Chaoyi Lu, Baojun Liu, Zhou Li, Shuang Hao, Hai-Xin Duan, Mingming Zhang, Chunying Leng, Ying Liu, Zaifeng Zhang, and Jianping Wu. An end-to-end, large-scale measurement of dns-over-encryption: How far have we come? In *Proceedings of the Internet Measurement Conference, IMC 2019, Amsterdam, The Netherlands, October 21-23, 2019*, pages 22–35. ACM, 2019.
- [30] Casey T. Deccio and Jacob Davis. DNS privacy in practice and preparation. In Aziz Mohaisen and Zhi-Li Zhang, editors, *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT 2019, Orlando, FL, USA, December 09-12, 2019*, pages 138–143. ACM, 2019.
- [31] Trinh Viet Doan, Irina Tsareva, and Vaibhav Bajpai. Measuring DNS over TLS from the edge: Adoption, reliability, and response times. In Oliver Hohlfeld, Andra Lutu, and Dave Levin, editors, *Passive and Active Measurement - 22nd International Conference, PAM 2021, Virtual Event, March 29 - April 1, 2021, Proceedings*, volume 12671 of *Lecture Notes in Computer Science*, pages 192–209. Springer, 2021.
- [32] Peter B. Danzig, Katia Obraczka, and Anant Kumar. An analysis of wide-area name server traffic: A study of the internet domain name system. In Deepinder P. Sidhu and David Oran, editors, *Proceedings of the Conference on Communications Architecture & Protocols, SIGCOMM 1992, Baltimore, Maryland, USA, August 17-20, 1992*, pages 281–292. ACM, 1992.
- [33] Jaeyeon Jung, Emil Sit, Hari Balakrishnan, and Robert Tappan Morris. DNS performance and the effectiveness of caching. In Vern Paxson, editor, *Proceedings of the 1st ACM SIGCOMM Internet Measurement Workshop, IMW 2001, San Francisco, California, USA, November 1-2, 2001*, pages 153–167. ACM, 2001.
- [34] Bernhard Ager, Wolfgang Mühlbauer, Georgios Smaragdakis, and Steve Uhlig. Comparing DNS resolvers in the wild. In Mark Allman, editor, *Proceedings of the 10th ACM SIGCOMM Internet Measurement Conference, IMC 2010, Melbourne, Australia - November 1-3, 2010*, pages 15–21. ACM, 2010.
- [35] John S. Otto, Mario A. Sánchez, John P. Rula, and Fabián E. Bustamante. Content delivery and the natural evolution of DNS: remote dns trends, performance issues and alternative solutions. In John W. Byers, Jim Kurose, Ratul Mahajan, and Alex C. Snoeren, editors, *Proceedings of the 12th ACM SIGCOMM Internet Measurement Conference, IMC '12, Boston, MA, USA, November 14-16, 2012*, pages 523–536. ACM, 2012.
- [36] Edith Cohen and Haim Kaplan. Proactive caching of DNS records: addressing a performance bottleneck. *Comput. Networks*, 41(6):707–726, 2003.
- [37] KyoungSoo Park, Vivek S. Pai, Larry L. Peterson, and Zhe Wang. Codns: Improving DNS performance and reliability via cooperative lookups. In Eric A. Brewer and Peter Chen, editors, *6th Symposium on Operating System Design and Implementation (OSDI 2004), San Francisco, California, USA, December 6-8, 2004*, pages 199–214. USENIX Association, 2004.
- [38] Trinh Viet Doan, Justus Fries, and Vaibhav Bajpai. Evaluating public DNS services in the wake of increasing centralization of DNS. In Zheng Yan, Gareth Tyson, and Dimitrios Koutsonikolas, editors, *IFIP Networking Conference, IFIP Networking 2021, Espoo and Helsinki, Finland, June 21-24, 2021*, pages 1–9. IEEE, 2021.
- [39] Pavlos Sermpezis, Lars Prehn, Sofia Kostoglou, Marcel Flores, Athena Vakali, and Emile Aben. Bias in internet measurement platforms. In *7th Network Traffic Measurement and Analysis Conference, TMA 2023, Naples, Italy, June 26-29, 2023*, pages 1–10. IEEE, 2023.
- [40] Gijs Van Den Broek, Roland Van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. DNSSEC meets real world: Dealing with unreachability caused by fragmentation. *IEEE Communications Magazine*, 2014.
- [41] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. *Computer Communications Review*, 45(3):35–42, 2015.
- [42] Thomas Holterbach, Cristel Pelsser, Randy Bush, and Laurent Vanbever. Quantifying interference between measurements on the RIPE atlas platform. In Kenjiro Cho, Kensuke Fukuda, Vivek S. Pai, and Neil Spring, editors, *Proceedings of the 2015 ACM Internet Measurement Conference, IMC 2015, Tokyo, Japan, October 28-30, 2015*, pages 437–443. ACM, 2015.
- [43] Vaibhav Bajpai et al. Vantage point selection for IPv6 measurements:

- Benefits and limitations of RIPE Atlas tags. In *IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017.
- [44] Pratyush Dikshit. Dotcp ripe dataset and probe resolvers. <https://zenodo.org/uploads/10632827>, 2024.
- [45] Pratyush Dikshit. Dns_resiliencybydotcp. https://github.com/pratyushdikshit26/DNS_ResiliencyByDoTCP, 2024.
- [46] Giovanni Schmid. Thirty years of DNS insecurity: Current issues and perspectives. *IEEE Commun. Surv. Tutorials*, 23(4):2429–2459, 2021.
- [47] Haya Schulmann. Pretty bad privacy: Pitfalls of DNS encryption. In Gail-Joon Ahn and Anupam Datta, editors, *Proceedings of the 13th Workshop on Privacy in the Electronic Society, WPES 2014, Scottsdale, AZ, USA, November 3, 2014*, pages 191–200. ACM, 2014.
- [48] Carlo Contavalli, Wilmer van der Gaast, David C. Lawrence, and Warren Kumari. Client subnet in DNS queries. *RFC*, 7871:1–30, 2016.
- [49] Jiarun Mao, Michael Rabinovich, and Kyle Schomp. Assessing Support for DNS-over-TCP in the Wild. In *Passive and Active Measurement Conference (PAM)*, pages 487–517. Springer, 2022.
- [50] Christian Huitema, Sara Dickinson, and Allison Mankin. DNS over dedicated QUIC connections. *RFC*, 9250:1–27, 2022.
- [51] Mike Kosek, Luca Schumann, Robin Marx, Trinh Viet Doan, and Vaibhav Bajpai. DNS Privacy with Speed?: Evaluating DNS over QUIC and its Impact on Web Performance. In *IMC*, pages 44–50. ACM, 2022.
- [52] Jayasree Sengupta, Mike Kosek, Justus Fries, Pratyush Dikshit, and Vaibhav Bajpai. Web Privacy By Design: Evaluating Cross-layer Interactions of QUIC, DNS and H/3. In *2023 IFIP Networking Conference, IFIP Networking*, 2023.
- [53] Jayasree Sengupta, Mike Kosek, Justus Fries, Simone Ferlin, Pratyush Dikshit, and Vaibhav Bajpai. On cross-layer interactions of quic, encrypted DNS and HTTP/3: design, evaluation and dataset. *CoRR*, abs/2306.11643, 2023.
- [54] Jana Iyengar and Martin Thomson. QUIC: A UDP-Based Multiplexed and Secure Transport. *RFC*, 9000:1–151, 2021.
- [55] Mike Kosek, Trinh Viet Doan, Malte Granderath, and Vaibhav Bajpai. One to Rule Them All? A First Look at DNS over QUIC. In *Passive and Active Measurement Conference*, 2022.
- [56] Sebastián García, Karel Hynek, Dmitrii Vekshin, Tomás Cejka, and Armin Wasicek. Large scale measurement on the adoption of encrypted DNS. *CoRR*, abs/2107.04436, 2021.
- [57] Danilo Cicalese, Diana Joubblatt, Dario Rossi, Marc-Olivier Buob, Jordan Augé, and Timur Friedman. A fistful of pings: Accurate and lightweight anycast enumeration and geolocation. In *2015 IEEE Conference on Computer Communications, INFOCOM 2015, Kowloon, Hong Kong, April 26 - May 1, 2015*, pages 2776–2784. IEEE, 2015.
- [58] Dan Komosny, Miroslav Voznák, and Saeed Ur Rehman. Location accuracy of commercial IP address geolocation databases. *Inf. Technol. Control.*, 46(3):333–344, 2017.
- [59] Ondrej Krajsa and Lucie Fojtova. RTT measurement and its dependence on the real geographical distance. In Norbert Herencsar and Karol Molnar, editors, *34th International Conference on Telecommunications and Signal Processing (TSP 2011), Budapest, Hungary, Aug. 18-20, 2011*, pages 231–234. IEEE, 2011.
- [60] Valeria Cardellini, Michele Colajanni, and Philip S. Yu. Geographic load balancing for scalable distributed web systems. In *MASCOTS 2000, Proceedings of the 8th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 29 August - 1 September 2000, San Francisco, California, USA*, pages 20–27. IEEE Computer Society, 2000.
- [61] Jian Zhang, T. Hamalainen, J. Joutsensalo, and K. Kaario. Qos-aware load balancing algorithm for globally distributed web systems. In *2001 International Conferences on Info-Tech and Info-Net. Proceedings (Cat. No.01EX479)*, volume 2, pages 60–65 vol.2, 2001.
- [62] Nandita Dukkipati, Tiziana Refice, Yuchung Cheng, Jerry Chu, Tom Herbert, Amit Agarwal, Arvind Jain, and Natalia Sutin. An argument for increasing tcp's initial congestion window. *Comput. Commun. Rev.*, 40(3):26–33, 2010.



Pratyush Dikshit is a Researcher at CISPA Helmholtz Center for Information Security, Germany. He received his M.S. (by Research) degree in Computer Science and Engineering specializing in cryptography from the National Institute of Technology, Tiruchirappalli, India (2017). Before joining CISPA in 2022, he worked for Cisco Systems as a Network Consulting Engineer for 5 years. His research interests include transport layer security, DNS privacy, Internet resiliency, and cryptography.



Mike Kosek is a Researcher at the Chair of Connected Mobility at the Technical University of Munich, Germany. His research focuses on Internet measurements in general, and transport protocol standardization, development, and deployment, in particular.



Nils Faulhaber received his M.Sc. from the Technical University of Munich in 2022 discovering his enthusiasm for Internet and IoT protocols as well as large scale data analysis. Subsequently, he joined Membrain where he is currently serving as Head of IoT and Analytics.



Jayasree Sengupta is currently working as a Post-doctoral Researcher at CISPA Helmholtz Center for Information Security, Germany. She received her PhD (2022) in Computer Science from the Indian Institute of Engineering Science and Technology, Shibpur, India, and MTech (2017) in Distributed and Mobile Computing from Jadavpur University, India. She is a recipient of the IFIP Networking Best Paper Award (2023) and JNCA Best Survey Paper Award (2022). Her research interests include Applied Cryptography, Blockchains, Network Security, and Data

Privacy.



intensive methods and by building real-world systems and models.

Vaibhav Bajpai is Professor and Head of the data-intensive Internet computing research group at the Hasso Plattner Institute and the University of Potsdam. Previously, he was an independent research group leader at the CISPA Helmholtz Center for Information Security, Hannover. Before that, he was a senior researcher at the Department of Computer Science at the Technical University of Munich. He received his PhD (2016) and Masters (2012) degrees from Jacobs University Bremen. His research focuses on improving Internet operations using data-intensive methods and by building real-world systems and models.