Evaluating Public DNS Services in the Wake of Increasing Centralization of DNS

Trinh Viet Doan, Justus Fries, Vaibhav Bajpai
Technical University of Munich, Germany
[trinhviet.doan | justus.fries | vaibhav.bajpai]@tum.de

Abstract—Recent studies have shown centralization in the Domain Name System (DNS) around public DNS services, which are hosted on centrally managed infrastructure and advertise higher reliability, improved security, and faster response times for name resolutions. However, many of the recently emerged public DNS services have not yet been extensively studied regarding popularity and performance. In light of this, we use 10.6k RIPE Atlas probes and find that 28.3% of the probes (and the their host network by extension) use at least one public DNS service, with Google being the most popular public DNS service among these probes. We further quantify the response time benefits of such public DNS services using ≈2.5k RIPE Atlas probes deployed in home networks (1k of which are IPv6 capable): Overall, we provision around 12.7M DNS requests based on a set of 23 domains and ten centralized public DNS services both over IPv4 and IPv6. For comparison, we additionally resolve the same set of domains using the probes' local resolvers, which are typically managed by the ISP and exhibit lower response times in general. We observe that even though IP and AS paths to local resolvers are generally shorter, some public DNS services (e.g., Cloudflare), achieve faster responses over both IPv4 and IPv6. Across all continents, Cloudflare, Google, and OpenDNS exhibit the lowest response times out of all public resolvers for successful DNS measurements. Probes in Europe (EU) and North America (NA) experience comparable latencies to public and local resolvers, thereby diminishing claimed latency benefits of public resolvers. We also observe inflated path lengths to and response times (over both address families) from most public resolvers for probes in Africa (AF) and South America (SA). Based on our observations, we provide recommendations and discuss situations in which switching to public DNS services may be beneficial.

I. INTRODUCTION

The Domain Name System (DNS) is said to become increasingly centralized [1], [2], concentrating around a small number of public DNS resolver services. These services are typically free of charge and promise increased reliability, faster response times, and higher security. In particular, as early supporters of the recently standardized DNS over TLS (DoT) [3], [4], [5], [6] and DNS over HTTPS (DoH) [7], [8] protocols [9], [10], [11], public DNS services such as Google, Cloudflare, Clean-Browsing, and Quad9 are gaining more traction and usage. Nevertheless, the latency differences between such centralized public resolvers and default ISP resolvers have not been extensively studied yet. Previous studies (§ II) have investigated the usage and performance of primarily two public DNS services: Google and OpenDNS. These studies found local ISP resolvers were more commonly used and provided better performance in

TABLE I
OVERVIEW OF THE PUBLIC DNS SERVICES MEASURED IN THE EXPERIMENT. ALL RESOLVERS THAT LAUNCHED AFTER 2010 HAVE NOT BEEN EXTENSIVELY STUDIED BEFORE.

Launch		IPv4 Address	IPv6 Address
2020-05	NextDNS	45.90.28.0	2a07:a8c0::
2018-04	Cloudflare DNS	1.1.1.1	2606:4700:4700::1111
2017-11	Quad9	9.9.9.9	2620:fe::9
2017-02	CleanBrowsing	185.228.168.168	2a0d:2a00:1::1
2017-02	Neustar UltraRecursive	156.154.70.1	2610:a1:1018::1
2015-09	VeriSign Public DNS	64.6.64.6	2620:74:1b::1:1
2013-11	Yandex DNS	77.88.8.8	2a02:6b8::feed:ff
2009-12	Google Public DNS	8.8.8.8	2001:4860:4860::8888
2006-07	OpenDNS	208.67.222.123	2620:0:ccc::2
2000-06	OpenNIC	185.121.177.177	2a05:dfc7:5::5353

terms of response times and proximity of the resolved location. Since publication of these studies (more than five years ago), an increasing number of new public DNS services have emerged (Table I), for which comparable studies are missing.

Due to the evolution of the DNS around such centralized public DNS services in recent years [12], we quantify the popularity, closeness w.r.t. path lengths, and performance benefits regarding terms of response times of *public resolvers* in comparison to *local resolvers* (assigned by the ISP) to provide a better understanding of these newly launched services. To this end, we use the RIPE Atlas platform [13], making use of 2,502 probes to issue and measure DNS lookups toward ten centralized public DNS services, along with lookups using each probe's locally configured default resolvers. We repeat these measurements daily for a set of 23 domains over both IPv4 and IPv6 for a period of two weeks (§ III). We further perform `traceroute` measurements from the probes toward the public resolvers as well as the publicly routable IP addresses of local resolvers. Our main findings are:

Popularity (§ IV) – We determine the popularity of public DNS services among all 10.6k connected RIPE Atlas probes and find that 28.3% (3k) of the probes use at least one public DNS service as their locally configured resolver. Further, 12.9% (1.4k) of the probes only use public resolvers, rather than resolvers managed by the ISP; in particular, 9.2% (1k) of all probes exclusively use one single public DNS service as their default resolver. Google provides the most prevalent public DNS service, used by 78.4% of these 3k probes.

Path Lengths (§ V) – As expected, IP paths to local resolvers of ISPs are shorter (1–12 IP hops) compared to public resolvers (5–17 IP hops). We see that over IPv4 around 82%

of the local resolvers are located in the AS of the probe, i.e., the first AS hop, with the number being even higher over IPv6 (93%). In contrast, AS paths to public resolvers involve around 2–5 ASes over both address families. Google Public DNS (80–86% samples, 2 AS hops) directly peers with the ISP, while Cloudflare and Quad9 (92–94% samples, 3 AS hops) tend to have an additional transit AS in between. We notice that Google edge caches deployed inside the ISP do not (yet) offer DNS services. We also observe that probes in South America (SA) exhibit higher IP and AS path lengths toward all DNS resolvers than any other continent.

Response Times (§ VI) – We find that 75% of all successful DNS requests are responded to within 40 ms over both address families. Unlike previous studies, we find that some public DNS services achieve lower lookup latency compared to local ISP resolvers over both address families. Responses from local resolvers are faster for 36–60% of the samples over IPv4 and 29–60% over IPv6, respectively. Specifically, probes in all continents besides Europe (EU) and North America (NA) experience worse response times from public resolvers, which shows overall benefits of local resolvers for substantial latency improvements over both address families (26.6 ms over IPv4, 51.8 ms over IPv6 on average). We also notice inflated response times to Google Public DNS for probes in Africa (AF), which indicates fewer points of presence in this continent. Yet, DNS response times for probes in AF and SA are significantly worse over IPv6 than IPv4, indicating the need to strengthen performance over IPv6 in these regions.

Based on these observations, we discuss (§ VII) recommendations, e.g., in which cases switching to public DNS services can provide performance benefits, along with limitations of the study. The data is publicly available via the RIPE Atlas API; we share the measurement IDs along with the analysis scripts and Jupyter notebooks to ease reproducibility of our work[1].

II. RELATED WORK

One of the first studies to measure performance of public DNS resolvers is presented by Ager *et al.* [14] (2010). They compare the responsiveness, the deployment, and the answers of local DNS resolvers to two public resolvers, namely Google DNS and OpenDNS. Performing active measurements using 60 vantage points in 28 different countries and 5 continents, they find that local resolvers managed by the ISPs generally outperform public resolvers in terms of response times for the most part. In addition, they find that these centralized resolvers lack local information about the requester (unlike the ISP), i.e., recursive resolution cannot consider location data and, thus, refers clients to suboptimal server locations as a result.

Regarding popularity of public DNS services, Otto *et al.* [15] (2012) show that usage of public DNS grows by 27% annually: Google is used by over 4% of the users, followed by OpenDNS with slightly over 3% as of 2011. They show that using public DNS services results in significantly different CDN redirections, often leading to degraded HTTP performance. To overcome this lack of additional information about the clients, Extended DNS (EDNS)[16] has been introduced, which allows clients to include an IP address prefix to the DNS request, so that the responses of recursive resolvers can be based on more informed decisions. Moreover, Callahan *et al.* [17] (2013) analyze 200M DNS queries passively measured in 90 homes in the USA with their corresponding 162M DNS responses. They find that Google's public DNS service is used in slightly over 1% of the queries, while 97% of the requests go to the ISP local resolver. Providing an additional view, APNIC Labs [18] (2014) use Javascript code embedded in advertisements to send DNS queries to a controlled authoritative DNS server. Using this technique, they show largely different numbers, as they identify 10.5% of the users to leverage Google's DNS service; these users are mostly located in Middle America, Central Africa, the Middle East, and South East Asia.

In recent years, several studies investigate DNS centralization from different points of view. For instance, Allman [19] (2018) analyzes the shared infrastructure w.r.t. Second-Level Domains (SLDs). He finds that Cloudflare and GoDaddy are the DNS providers with by far the highest numbers of SLDs managed, each accounting for roughly 70k SLDs outsourced to them, in comparison with the remaining providers in the top 10 (which amount to 204k SLDs in total). Similarly, Zembruzki *et al.* [20] (2020) develop `dnstracker`, an active measurement tool that enables the assessment of the levels of concentration and shared infrastructure in the DNS. The tool first resolves a domain via `dig` to learn about the associated authoritative name server, before it runs `traceroute` to measure the path to the authoritative server. In the recorded trace, the Hop-Before-The-Last (HBTL) then indicates the AS and hosting DNS provider. They measure the Alexa Top 1M domains' authoritative name servers with their tool, finding that up to 12k name servers share the same infrastructure, which may result in single points of failure. Moura *et al.* [21] (2020) measure DNS traffic at a DNS root server and two Top-Level Domains (TLDs) (`.nl` and `.nz`), i.e., the traffic between recursive resolvers and authoritative servers. In particular, they focus on five cloud and CDN providers, namely Google, Amazon, Microsoft, Facebook, and Cloudflare. They find that the centralization benefits the deployment of DNS features such as DNSSEC or QNAME minimization, along the usage of IPv6, as these big players push these features, although the adoption varies between providers. Further, they show that DNS traffic is centralized around the five providers: For the root server, around 9% of the traffic is received from those providers, whereas for the country-code TLDs, they find more than 30% of the incoming queries to originate from the five providers. Overall, these studies show moderate centralization in the DNS from different perspectives.

The trend of DNS centralization can also be seen outside of empirical studies. For instance, Mozilla aimed to enable DoH using Cloudflare's resolver by default for all Firefox users, though this proposal received pushback due to lack of control and privacy for users [22]. In response, Mozilla added an alternative trusted DoH resolver (NextDNS), starting rollout

---

[1] https://github.com/tv-doan/ifip-net-2021-dns

Fig. 1. Map of the geographical locations of the roughly 2.5k RIPE Atlas probes used in the measurement experiment.

in the US in February 2020 [23]. Similarly, Google Chrome auto-upgrades all its users (starting with Chrome 83) to DoH that use a DNS service that supports HTTPS connections [24]. However, DNS data is considered highly sensitive since it allows tracking of user behavior [25], [26], [27], [28], [29], meaning that public resolvers can pose a risk to privacy [1], [30] as a result of consolidation and centralization in the DNS [2], [31], [32]. Related to these trends, note that Neustar acquired VeriSign Public DNS in November 2020 [33] (shortly after our measurement experiment).

As can be seen, with the advent of centralized public DNS services offered by large organizations [12] and increased IPv6 adoption, the scenery of DNS operations has evolved over the years. Given that these services have not been extensively studied in recent years, in particular regarding path lengths and latency toward them and from the perspective of home users, this paper attempts to provide a broader perspective on the popularity and performance benefits of such centralized DNS services in this changed landscape.

III. METHODOLOGY

We use the RIPE Atlas platform [13] to perform DNS lookup and `traceroute` measurements for a variety of DNS resolvers.

Measurement Probes: As older RIPE Atlas hardware probes (versions 1 and 2) are known to be affected by load issues, we choose hardware probes that are tagged as at least version 3 (V3), since these probes are less sensitive to load [34], [35]. Additionally, we pick home probes that have native IPv4 and/or native IPv6 connectivity, using RIPE Atlas tags [36] for the selection. RIPE Atlas anchor probes are not considered as we are mainly interested in DNS resolution for end users. In this way, we leverage 2,502 probes hosted in 729 distinct Autonomous Systems (ASes) across 89 countries to perform our measurements (see Fig. 1). Out of these 2.5k home probes, 2,491 probes (99.6%) are IPv4-capable, 1,090 probes (43.6%) IPv6-capable.

Probes may use public DNS services (see Table I) as their locally configured (default) DNS resolver, which they use for "on-probe" name resolutions. We exclude such DNS measurements (i.e., on-probe but toward public resolvers) from our analyses and will use the term *local resolvers* for the remaining cases, i.e., measurements to resolver endpoints that

are *not* assigned to one of the public DNS services. For the identification of such cases among locally configured resolvers, we also consider alternative IP addresses of the public services, e.g., including `8.8.4.4` in addition to `8.8.8.8` for Google (using an extended list which we will share along with the other analysis artifacts). These alternative endpoints are, for instance, used by the services for load balancing purposes or to provide specific filters for unwanted domains. Note that in case a RIPE Atlas probe is provided with multiple IP addresses for the locally configured resolvers, e.g., with one ISP, one Google, and one Cloudflare resolver endpoint, creating one on-probe DNS measurement via RIPE Atlas will cause the probe to issue queries toward *all* these resolvers.

DNS Queries (§ VI): Over a period of two weeks in September 2020 and from each probe, we issue daily DNS lookups over UDP/53 for a set of domains toward ten selected public resolvers (see Table I). We also issue the same lookups for these domains to the probe's local DNS resolvers (managed by the ISP and assigned via DHCP unless explicitly reconfigured by the probe host). Queries are sent over both address families: We query `A` records over IPv4 and `AAAA` records over IPv6. We choose a subset of dual-stacked domains from Alexa Top 1M [37], where some website domains are served by Content Delivery Networks (CDNs) (20 domains), while the remaining ones (2 domains) are non-CDN hosted website domains. However, note that we do not find significant response time differences for the different domains, which is why we do not separate the analysis by Alexa rank or by CDNs used: The repeated DNS queries toward the same domains and their popularity based on the Alexa toplist provide an increased probability of all DNS records being cached. We further leave the Recursion Desired (RD) flag unset for measurements toward the public resolvers, which nullifies other potential latency differences regarding recursive lookups; note that DNS requests using the configured on-probe resolvers will always have the RD flag set (due to RIPE Atlas policies for privacy reasons). However, since the records are likely cached as discussed above, recursive lookups are unlikely to occur for local resolvers as well.

In addition to the 22 domains described above, we issue queries for a set of fabricated (nonexistent) domains, namely `$r.google.com`, where `$r` is a random 16-digit hex string created by a probe for each measurement run. As this domain name is virtually guaranteed to be unique and not cached as a result, resolvers should ultimately return `NXDOMAIN` messages to the query. Although DNS wildcards [38] could cause a resolver to return a non-`NXDOMAIN` message instead, a wildcard for the chosen SLD is not likely to exist. When analyzing the fabricated domains, we find that some resolvers do not return `NXDOMAIN` messages to the randomized domain queries; e.g., the local resolvers of some IPv6-capable AT&T probes return an IP address that redirects to `dnserrorassist.att.net`, a Web service that performs a Web search for the nonexistent domain using a search engine. Such cases indicate hijacking of DNS responses [39] by the ISP, given they should instead

return `NXDOMAIN` for nonexistent and fabricated domains.

Overall, we collect data for around 12.7M DNS requests (14 days × [2,491 (IPv4) + 1,090 (IPv6) probes] × 23 domains × at least 11 resolvers) in cooperation with RIPE NCC. This results in a set of 506 measurement IDs (2 address families × 23 domains × 11 target resolvers), with each ID grouping the recurring measurements for the selected probes. We are aware of methodologies that include DNS records of an authoritative server controlled by the researchers [40]. However, we value such controlled experiments as separate investigations and consider them orthogonal to the goal of this paper, given our experiment is designed to resemble the perspective of end users that do not necessarily have such control knobs.

Traceroute (§ V): We additionally run one-off ICMP `traceroute` measurements toward the resolver endpoints (Table I). The measurements are performed from each probe over IPv4 and/or IPv6. Since RIPE Atlas does not allow `traceroute` measurements toward private IP addresses [41], only local resolvers with public IP addresses can be traced: We determine the addresses of all public local resolvers and additionally run `traceroute` toward these endpoints (§ V). As centralized public resolvers leverage IP anycast [42], we expect IP paths from the probe to these resolvers to not be inflated unnecessarily. Thus, inflated path lengths would reveal cases in which centralized public resolvers lack points of presence.

Additional On-Probe Resolutions (§ IV): To estimate the popularity of centralized public DNS resolvers among RIPE Atlas probes, we additionally take all connected probes (10.6k probes) into account and issue the domain `google.com` to be resolved *on the probe*, i.e., via the locally configured resolvers. The source IP addresses stated in the DNS responses allow us to determine whether (and how many) public DNS services and/or local ISP resolvers are used by the probes for name resolution. Note that repeating the query for each on-probe resolver is not necessary, since RIPE Atlas probes will automatically query the domain using *all* resolvers listed in their DNS configuration.

IV. RIPE ATLAS PROBE RESOLVER BIAS

We begin the analysis by investigating the popularity of centralized public DNS services on the RIPE Atlas measurement platform. Using on-probe resolutions for all 10,624 connected probes, we study how biased on-probe resolutions are toward public resolvers. We find that 7,617 probes employ local (i.e., non-centralized) resolvers exclusively for their DNS lookups (71.7%), while the remaining 3,007 probes use at least one public DNS service for on-probe name resolution (28.3%).

Table II provides an overview of the usage of public services among the latter 3k probes. In particular, 1,636 probes leverage a combination of local and public resolvers (15.4%), whereas 1,371 probes use public resolvers exclusively as their default resolvers (12.9%). For the probes that leverage public resolvers exclusively, we observe that 978 probes only use a single service (71.3%). On the other hand, 355 probes use two different services (25.9%), while 38 probes even make use of three (2.8%). In the first case, 978 probes rely on a single public DNS service for name resolution, which poses a privacy risk

	# Probes	# Probes with n Publ. Services	# Employing Probes
Public only	1,371 (12.9%)	978, $n = 1$ (71.3%) 355, $n = 2$ (25.9%) 38, $n = 3$ (2.8%)	Google: 1,001 (55.5%) Cloudflare: 527 (29.2%) Quad9: 126 (7.0%) OpenDNS: 122 (6.8%) Yandex: 12 (0.7%) NextDNS: 8 (0.4%) VeriSign: 3 (0.2%) Neustar: 2 (0.1%) CleanBrowsing: 1 (<0.1%)
Public + local	1,636 (15.4%)	825, $n = 1$ (50.4%) 811, $n = 2$ (49.6%)	Google: 1,357 (56.7%) VeriSign: 656 (27.4%) Cloudflare: 263 (11.0%) OpenDNS: 54 (2.3%) Quad9: 47 (2.0%) Yandex: 13 (0.5%) Neustar: 2 (0.1%) NextDNS: 2 (0.1%) OpenNIC: 1 (<0.1%)

(also for the users of the home network by extension), as the entire outbound DNS traffic gets consolidated to a single third-party that has to be fully trusted [43] in addition to the ISP (which handles the user's Internet traffic in general).

Among the 1.4k probes that use public resolvers exclusively, we determine Google to be the most popular service: Google is used in 1,001 samples (55.5%) of the observed probe-resolver pairs (i.e., 1×978 + 2×355 + 3×38 = 1,802 unique pairs of probes and resolvers) and shows a higher prevalence than all other services combined as a result (Cloudflare 29.2%, Quad9 7.0%, OpenDNS 6.8%, Yandex 0.7%, NextDNS 0.4%, VeriSign 0.2%, Neustar 0.1%, CleanBrowsing <0.1%).

Regarding the probes that use both local and public resolvers, we find that around half of the probes make use of one public DNS service in addition to their local resolvers (825 probes, i.e., 50.4%); the other half (811 probes) uses even two additional public DNS services (49.6%). We see that Google is also the most popular public DNS service (56.7%) to complement local resolvers. We find that VeriSign is the second most common service with 27.4%, followed by Cloudflare (11.0%), OpenDNS (2.3%), Quad9 (2.0%), and Yandex (0.5%), while Neustar, NextDNS, and OpenNIC account for ≤0.1% each. As such, we observe that VeriSign is a much more popular DNS service among RIPE Atlas probes when used as a complement for local resolvers as opposed to a standalone service.

In conclusion, roughly 3k (28.3%) of all 10.6k RIPE Atlas probes use at least one public resolver for on-probe resolution, with 1,204 probes (11.3%) leveraging at least two different public services; Google is the most prevalent public DNS service, used by 2,358 of the 3k probes (78.4%). Measurement studies should take this into account (§ VII) when performing DNS measurements via RIPE Atlas and selecting to resolve the names "locally" on the probe, since results may include responses from public resolvers as well and could, therefore, lead to unintended side-effects. Moreover, as mentioned before,

Google Chrome users (starting with Google Chrome 83) are auto-upgraded to DoH [24] when using at least one of the public DNS services from CleanBrowsing, Cloudflare, Google, NextDNS, OpenDNS, or Quad9 (among others). Our dataset reveals that this upgrade policy would affect at least 2,991 (28.2%) of the users hosting RIPE Atlas probes if they used Google Chrome for browsing the Web.

V. PATH LENGTHS

We collect 32k successful `traceroute` measurements from the 2.5k probes toward the ten centralized public DNS services and the publicly routable local resolvers over IPv4 (2.5k probes) and/or IPv6 (1k probes). Fig. 2 shows the measured IP (top) and AS (bottom) path lengths; in all plots, the number of successful samples is specified in the legends for each resolver, i.e., (#IPv4, #IPv6) samples.

IP Path Lengths: We find that IP path lengths to public resolvers are between 5–17 IP hops over both address families for the nearly all samples. Local resolvers (that are publicly routable) tend to be closer to the probes, as almost all IP paths are only 1–12 IP hops long. While IPv6 paths are marginally shorter than IPv4 paths for most samples, larger differences of 2–3 hops are only visible for Google. When comparing the IP path lengths between continents (plots not shown), we see less pronounced differences between IPv4 and IPv6 within a continent. However, probes in SA and Oceania (OC) experience moderately higher IP path lengths toward all DNS resolvers than the other continents, which are by and large comparable regarding IP path lengths.

AS Path Lengths: We further determine the AS paths to the resolvers based on the IP path measurements. We lookup the ASes of the encountered IP address prefixes via RIPEstat, which derives AS information from BGP data collected by RIPE's Routing Information Service (RIS). The first IP hop represents AS hop 1; whenever the AS announcing the IP prefix of the intermediate IP hop changes, we increment the AS hop by 1. In case routers along the path are non-responsive or do not have announcing ASes based on RIPEstat, we keep the current AS hop count, but drop measurements with more than two missing AS mappings to avoid incorrect counting.

We observe that nearly 82% of the `traceroute` measurements toward the local resolvers over IPv4 end within the same AS they originated from, i.e., AS path lengths of 1. Over IPv6, 93.7% of the samples end in the originating AS. Overall, nearly all AS paths to local resolvers involve at most 2–3 ASes with a maximum of 7 ASes for IPv4 and 6 ASes for IPv6.

In comparison, public DNS resolvers are farther away, although they are still relatively close to the probes: 80.4% of the samples toward Google over IPv4 (85.8% over IPv6) are located in the AS after the origin AS, i.e., have AS path lengths of 2, meaning that Google directly peers with ISPs. On the other hand, this means that Google edge caches (deployed in ISP networks) do not yet offer DNS services toward clients; enabling this could be one approach to reducing the latency to Google Public DNS even further in the future. Also, we see traces toward Cloudflare and Quad9 to be similarly long
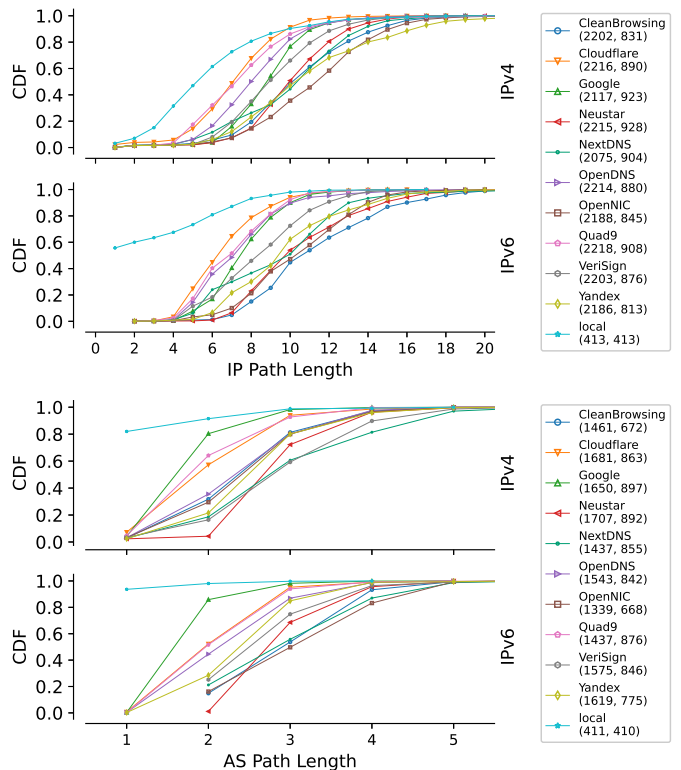

Fig. 2. Distribution of IP (top) and AS (bottom) path lengths. Most local resolvers are located in the probe's AS and have the shortest IP and AS path lengths. In contrast, most public resolvers have AS path lengths of 2–3.

as toward Google, although slightly longer, as 57.3% and 64.2% respectively of each of the samples exhibit an IPv4 AS path length of 2, with around 94% and 92.7% at an AS path length of 3 over IPv4 (52.2% and 51.7% for AS path length of 2, and 95.3% and 94% for AS path length of 3 over IPv6, respectively). The remaining public resolvers have higher AS path lengths, although most samples have AS path lengths of 2–5 over IPv4 and IPv6, meaning that multiple transit ASes need to be crossed to reach those public DNS resolvers. Similar to IP paths, AS paths are also inflated in SA over both address families (roughly 40% of them longer than 3 AS hops). While these observations suggest both flattening of the Internet topology [44] and that public DNS resolvers are relatively close to the probes as a result, the measurements also reveal that resolvers, particularly in SA, could be moved closer to the edge.

VI. DNS RESPONSE TIMES

In order to quantify the latency to the measured resolvers, we analyze and compare the DNS response times for all successful samples, i.e., the time it takes from sending an initial DNS request until the DNS response arrives at the probe.

We specifically request `A` records over IPv4 and `AAAA` records over IPv6. However, some resolvers return additional resource records as well: For instance, both `A` and `AAAA` records are returned to a single `A` query over IPv4 (and `AAAA` query over IPv6, respectively). Note that RIPE Atlas reports individual

Fig. 3. Distribution of response times for each resolver. Local resolvers show higher response times than public resolvers for around 55% of the samples over IPv4; over IPv6, public resolvers generally respond faster in comparison.

results for each returned record type that were returned to the same request, including different response times for each of the record types. We notice that this behavior is exclusive to measurements that use the configured on-probe resolvers (regardless of whether a local or public resolver is queried). We find that most of these cases with mixed up record types originate from local resolvers (81.4% IPv4, 90.1% IPv6), although Google and Cloudflare also account for 8.7% and 7.0% of the samples over IPv4 (4.1% and 4.3% over IPv6, respectively). In rare cases (<1.5%), we also see this for Quad9, OpenDNS, Neustar, Yandex, and NextDNS. When comparing the response time differences between such responses, around 74% of the respective samples over IPv4 exhibit a difference of at most ± 10 ms between the A and AAAA records (64% over IPv6). Hence, response times between A and AAAA records requested over either address family are largely similar in most cases, which indicates that both A and AAAA records are cached and that the resolvers likely operate in dual-stack. Thus, we do not further distinguish between different record types and focus on differences over IPv4 and IPv6 instead. As mentioned in § III, we exclude on-probe DNS measurements that leverage public resolvers in the following and further only include measurements with a NOERROR response code and a non-empty answer section, which we consider as successful measurements.

Overall Distribution: Over IPv4, the interquartile range (IQR, i.e., 25^{th}–75^{th} percentiles) of the response times across all resolvers and probes is $[9.5; 36.8]$ ms, while 90.7% of the queries are responded to within 100 ms. Over IPv6, the IQR covers $[10.1; 39.9]$ ms instead; similarly, 86.9% of the responses are received within 100 ms over IPv6 as well, which suggests that both address families show comparable DNS response times overall.

By Resolver and Address Family: We further investigate the response times of different resolvers as observed by the probes. Fig. 3 shows the distributions of the response times of all successful DNS lookups for each resolver over IPv4 and IPv6. Note that while we observe Neustar and OpenNIC to

also exhibit low response times in general, the measurements are not quite comparable to other resolvers due to the much lower sample size of successful measurements, caused by high numbers of REFUSED responses. Consequently, we do not discuss Neustar and OpenNIC in detail in the following.

Generally, we observe Cloudflare, Google, and OpenDNS to achieve the lowest response times out of the public resolvers. NextDNS, Quad9, and CleanBrowsing show relatively similar performances (slightly behind the previous ones), whereas VeriSign and Yandex exhibit visibly higher response times.

1) IPv4: Over IPv4 (Fig. 3 top), most resolvers exhibit similar response time distributions, although we see some resolvers performing differently than others. For instance, around 40% of the samples for local resolvers have a DNS lookup time of ≤ 10 ms, with roughly 83.5% of the samples taking up to 100 ms. In comparison, each of the public resolvers responds within 10 ms for less than 20% of their successful responses only, except for Cloudflare (34.6%) and OpenDNS (25.2%), with Google (19.8%) slightly below 20%. At the 40^{th} percentile of local resolvers, only Cloudflare achieves similar response times with 11.2 ms; for the remaining 60% of the samples, local resolvers perform comparably to or worse than public resolvers. Thus, local resolvers exhibit more varying results in comparison with most public resolvers. Further, Yandex performs visibly worse than any other public resolver (IQR of $[40.2; 127.6]$ ms), however, note that Yandex primarily operates in Russia; this is also reflected in its inflated path lengths for all non-EU probes.

2) IPv6: The measurements over IPv6 (Fig. 3 bottom) show similar distributions and differences between resolvers as over IPv4. Over IPv6, we observe that only around 33.4% of the local resolver responses return within 10 ms, although response times of up to 100 ms are still achieved by 83.7% of the samples (83.5% over IPv4). Cloudflare responds to requests in less than 40 ms for 93.5% of the cases over IPv6, making it one of the fastest public resolvers, followed by Google (90.6%) and OpenDNS (82.7%). Moreover, we observe that response times for CleanBrowsing are substantially shifted toward the higher end for around 40% of the samples, which suggests that its IPv6 performance lags behind in some scenarios.

By Continent: Given the previous observations, we dig deeper by also taking regional distribution into account. Fig. 4 depicts the median response times w.r.t. median response times of probe and resolver pairs, aggregated by probe location (continent) for IPv4 (top) and IPv6 (bottom). Blank cells represent areas in which probes could not resolve any domain successfully. Note that the shown response times are median values of the median response times by probe and, thus, not directly comparable to Fig. 3. Moreover, recall that the number of samples for Neustar and OpenNIC are much lower and distributed differently compared to the other resolvers (see above), which explains their unusual values seen in Fig. 4.

Over IPv4, we observe that probes in EU and NA exhibit much lower response times (18.4 ms and 29.6 ms on average) compared to other continents (where the averages range from 50.6 ms up to 90.8 ms), particularly when compared to probes
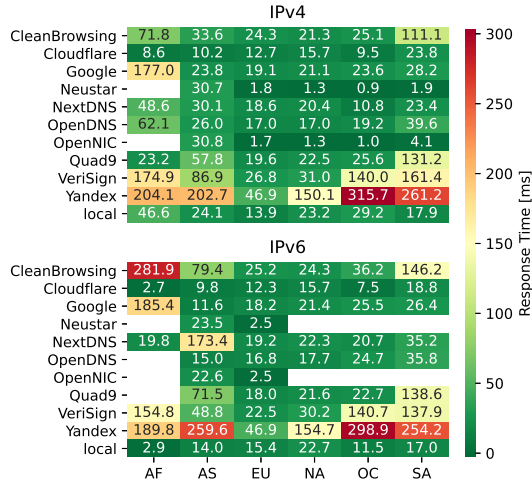
Fig. 4. Median DNS response times for each (probe, resolver) medians by continent. Probes in EU and NA measure lower response times than in SA and AF. Response times over IPv6 are more varying than over IPv4.

IPv4

	AF	AS	EU	NA	OC	SA
CleanBrowsing	71.8	33.6	24.3	21.3	25.1	111.1
Cloudflare	8.6	10.2	12.7	15.7	9.5	23.8
Google	177.0	23.8	19.1	21.1	23.6	28.2
Neustar		30.7	1.8	1.3	0.9	1.9
NextDNS	48.6	30.1	18.6	20.4	10.8	23.4
OpenDNS	62.1	26.0	17.0	17.0	19.2	39.6
OpenNIC		30.8	1.7	1.3	1.0	4.1
Quad9	23.2	57.8	19.6	22.5	25.6	131.2
VeriSign	174.9	86.9	26.8	31.0	140.0	161.4
Yandex	204.1	202.7	46.9	150.1	315.7	261.2
local	46.6	24.1	13.9	23.2	29.2	17.9

IPv6

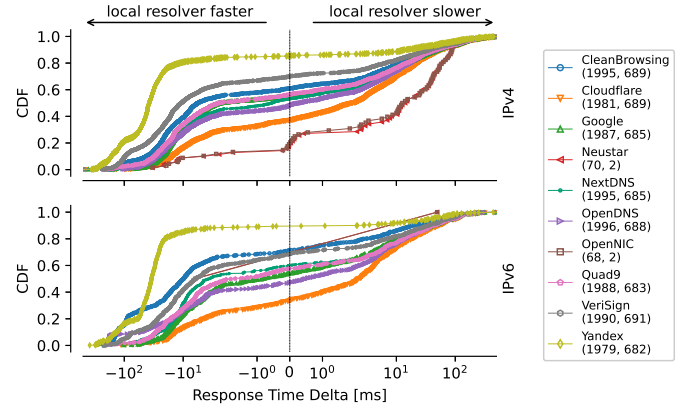|	AF	AS	EU	NA	OC	SA
CleanBrowsing	281.9	79.4	25.2	24.3	36.2	146.2
Cloudflare	2.7	9.8	12.3	15.7	7.5	18.8
Google	185.4	11.6	18.2	21.4	25.5	26.4
Neustar		23.5	2.5			
NextDNS	19.8	173.4	19.2	22.3	20.7	35.2
OpenDNS		15.0	16.8	17.7	24.7	35.8
OpenNIC		22.6	2.5			
Quad9		71.5	18.0	21.6	22.7	138.6
VeriSign	154.8	48.8	22.5	30.2	140.7	137.9
Yandex	189.8	259.6	46.9	154.7	298.9	254.2
local	2.9	14.0	15.4	22.7	11.5	17.0

Fig. 5. Distribution of median response time differences (between local and public resolvers) for each probe and public DNS service. Cloudflare and OpenDNS perform better than local resolvers for more than half of the samples over both address families.

in SA (73.1 ms) and AF (90.8 ms). In all continents besides AF, local resolvers achieve lower response times compared to public resolvers. However, for probes in OC, using either Cloudflare, NextDNS, or OpenDNS would result in improvements of 10–20 ms (see Fig. 4). The higher response times seen for Google (177.0 ms) from probes in AF indicates fewer points of presence in this continent. Over IPv6, response times are more varying than over IPv4. Probes located in EU (average 18.2 ms) and NA (36.7 ms) experience the lowest response times similar to IPv4; probes in OC (65.4 ms) and Asia (AS) (66.3 ms) see moderately high response times for most resolvers. Measurements from AF (119.6 ms) and SA (90.0 ms) have the highest response times over IPv6, with some resolvers not responding successfully in these regions at all.

By Probe: We use the previously determined median response times for pairs of probe and resolver to calculate the differences between the local resolver and each public DNS service, as shown in Fig. 5. We observe that over IPv4, local resolvers are faster than public DNS services for roughly 36–60% of the probes for most resolvers (Cloudflare 36.0%, OpenDNS 46.5%, NextDNS 51.7%, Google 53.7%, Quad9 54.7%, CleanBrowsing 59.4%), and up to around 68–83% of the probes for the remaining ones (Yandex 82.6%, VeriSign 68.0%). The results over IPv6 are similar for most resolvers: Around 29–60% of the probes experience slower responses for DNS requests when querying most DNS services (Cloudflare 29.2%, OpenDNS 40.5%, Google 46.4%, Quad9 49.4%, NextDNS 51.6%, VeriSign 59.3%, CleanBrowsing 61.7%) compared to local resolvers. This indicates that Cloudflare and OpenDNS are faster than the local DNS resolvers of more than 50% of the probes over both address families.

For each probe, we calculate the average difference of its data points in Fig. 5 over each IPv4 and IPv6. Across all probes, the overall average difference is −0.3 ms over IPv4 and −6.3 ms over IPv6, i.e., local resolvers are faster. Excluding probes

from EU and NA from this increases the average difference to −26.6 ms and −51.8 ms, which suggests that in regions outside of EU and NA probes benefit substantially when using local resolvers over public DNS services. Thus, public resolvers in EU and NA exhibit similar response times when compared to local resolvers, which indicates that making a switch to public DNS services on average is not necessarily beneficial.

VII. DISCUSSION

Recommendations: Our experiments identify that it is not possible to determine the preferential order of resolvers, i.e., preferred main resolver and potential backup resolver(s) for the RIPE Atlas probes. Probes issue DNS requests to *all* resolvers in the list (with a maximum of three resolvers as observed in our study), without retaining their order. We recommend RIPE Atlas and/or probe hosts to explicitly tag the order of resolvers in the DNS measurements. We also recommend RIPE Atlas probe hosts that use one public DNS resolver exclusively (§ IV) to add further DNS services to not consolidate their DNS traffic to a single public DNS service. Further, we recommend measurement studies via RIPE Atlas to take the DNS service used by the probes into account, as we find many probes to use public DNS service exclusively, and the performance differs based on whether a local or a public DNS service is set as default on-probe resolver; in particular, this could lead to unintended side-effects. To circumvent this, probe hosts could add a tag that specifies the usage of a public resolver as one of the configured on-probe resolvers. The DNS response times (§ VI) indicate that users do not benefit substantially from switching to public DNS resolver services in the average case and may consider to keep using their local resolvers to avoid sharing sensitive outbound DNS traffic with large (CDN) providers. In specific cases, the dataset suggests that users can experience latency benefits and, consequently, better user experience when switching to certain public resolvers, in particular Google or Cloudflare. Users in AF may consider Cloudflare DNS though, given the DNS response times to Google Public DNS are still

higher due to fewer points of presence in this region. However, users should also consider their individual trust relationships and tradeoffs between privacy and latency before switching. We find considerably high latencies over IPv6 for users in AF and SA for many resolvers and recommend providers of public DNS services to strengthen peering in these regions.

Limitations and Future Work: The collected dataset inherits a geographical bias of RIPE Atlas probe deployment, given the number of probes in some regions are fairly limited. Therefore, note that the measurement results are not necessarily generalizable to the whole Internet. In particular, the popularity of public DNS services is biased by the population of RIPE Atlas probe hosts (who typically have networking experience), which means that the prevalence and popularity is likely not fully representative of the general population. Further, the RIPE Atlas API not allowing `traceroute` measurements toward private IP address ranges prevents further distinction and limits analyses of local resolvers in this study; note that ICMP `traceroute` packets may be treated differently by middleboxes than the UDP-based DNS traffic as well. ISP resolvers can also forward a query to public DNS resolvers, however, characterizing such indirect use [45] is left for future work.

We acknowledge that DNS response times are only one piece of a transmission sequence: The IP address returned for a lookup, which we do not consider in our study, can point a client to a closer or more distant endpoint, which impacts the overall latency (see § II). However, as this is not the focus of this study, we plan to consider the quality of responses regarding server selection as future work, especially in the context of anycast regarding both the resolver and the resolved IP endpoint. Considering we only analyze successful DNS measurements in our study, the evaluation of non-successful results can additionally reveal regions and resolvers with high failure rates or cases of DNS filtering and censorship (similar to the DNS hijacking for `NXDOMAIN` responses discussed above). We plan to perform and include additional measurements and analyses to also investigate recursive lookup behavior for uncached records and specific cases with exceptionally low or high response times, among other open questions.

VIII. CONCLUSION

Using DNS measurements from the RIPE Atlas platform, we determine the usage of public DNS services and find that around a quarter (3k) of all 10.6k probes incorporate centralized public DNS services for name resolution by default. We observe Google to be by far the most prevalent public DNS service configured in home networks, as it is configured in 78.4% of these probes. We perform a set of measurements toward local resolvers as well as ten centralized public DNS services over both IPv4 and IPv6 across a two week time period to analyze the performance of DNS resolvers. In addition, we also run `traceroute` measurements toward the resolver endpoints to determine IP and AS path lengths to the resolvers. We observe that local resolvers are closer to the probes, as expected, although public resolvers are also only either one or two AS hops farther away from the probes, which suggests centralized

services moving closer to the edge as a result of Internet flattening. However, some paths in specific regions (mostly SA) are inflated. Generally, centralized public DNS services, in particular Google and Cloudflare, provide lower response times over both address families. Nevertheless, local resolvers are similarly fast, diminishing benefits of making a switch to centralized public DNS services. In regions besides EU and NA, using local resolvers can offer substantial latency benefits (26.6 ms for IPv4 and 51.8 ms for IPv6 on average). In light of observations from previous work, we find this result intriguing; a causal reasoning requires further investigation.

Due to recent concerns about consolidation of Internet services such as DNS, this paper adds to the understanding of public DNS services by quantifying their usage and benefits. While local resolvers still account for the majority and provide comparable performance to faster public resolvers (at least in EU and NA), this might change in the future: In particular, the increasing adoption of DNS over TLS and DNS over HTTPS [10] (mainly pushed by centralized public DNS services) contributes to increasing DNS centralization, which poses many questions for future work. To facilitate the exploration of these open questions, we share the measurement IDs and analysis code with the community.

REFERENCES

[1] J. Livingood, M. Antonakakis, B. Sleigh, and A. Winfield, "Centralized DNS over HTTPS (DoH) Implementation Issues and Risks," Tech. Rep., Sep. 2019. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-livingood-doh-implementation-risks-issues-04

[2] G. Huston, "DNS resolver centrality," Sep. 2019, accessed 2021-Jan-21. [Online]. Available: https://blog.apnic.net/2019/09/23/dns-resolver-centrality/

[3] Z. Hu, L. Zhu, J. S. Heidemann, A. Mankin, D. Wessels, and P. E. Hoffman, "Specification for DNS over Transport Layer Security (TLS)," *RFC*, vol. 7858, 2016. [Online]. Available: https://doi.org/10.17487/RFC7858

[4] S. Dickinson, D. K. Gillmor, and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS," *RFC*, vol. 8310, pp. 1–27, 2018. [Online]. Available: https://doi.org/10.17487/RFC8310

[5] R. Houser, Z. Li, C. Cotton, and H. Wang, "An Investigation on Information Leakage of DNS over TLS," in *Conference on Emerging Networking Experiments And Technologies*. ACM, 2019, pp. 123–137. [Online]. Available: https://doi.org/10.1145/3359989.3365429

[6] T. V. Doan, I. Tsareva, and V. Bajpai, "Measuring DNS over TLS from the Edge: Adoption, Reliability, and Response Times," in *Passive and Active Measurement Conference*, vol. 12671. Springer, 2021, pp. 192–209. [Online]. Available: https://doi.org/10.1007/978-3-030-72582-2_12

[7] P. E. Hoffman and P. McManus, "DNS queries over HTTPS (doh)," *RFC*, vol. 8484, pp. 1–21, 2018. [Online]. Available: https://doi.org/10.17487/RFC8484

[8] T. Böttger, F. Cuadrado, G. Antichi, E. L. Fernandes, G. Tyson, I. Castro, and S. Uhlig, "An Empirical Study of the Cost of DNS-over-HTTPS," in *Internet Measurement Conference*. ACM, 2019, pp. 15–21. [Online]. Available: https://doi.org/10.1145/3355369.3355575

[9] C. Lu, B. Liu, Z. Li, S. Hao, H. Duan, M. Zhang, C. Leng, Y. Liu, Z. Zhang, and J. Wu, "An End-to-End, Large-Scale Measurement of DNS-over-Encryption: How Far Have We Come?" in *Internet*

Measurement Conference. ACM, 2019, pp. 22–35. [Online]. Available: https://doi.org/10.1145/3355369.3355580

[10] C. T. Deccio and J. Davis, "DNS Privacy in Practice and Preparation," in *Conference on Emerging Networking Experiments And Technologies*, 2019, pp. 138–143. [Online]. Available: https://doi.org/10.1145/3359989.3365435

[11] A. Hounsel, K. Borgolte, P. Schmitt, J. Holland, and N. Feamster, "Comparing the Effects of DNS, DoT, and DoH on Web Performance," in *The Web Conference*. ACM / IW3C2, 2020, pp. 562–572. [Online]. Available: https://doi.org/10.1145/3366423.3380139

[12] R. Radu and M. Hausding, "Consolidation in the DNS resolver market – how much, how fast, how dangerous?" *Journal of Cyber Policy*, vol. 5, no. 1, pp. 46–64, 2020. [Online]. Available: https://doi.org/10.1080/23738871.2020.1722191

[13] RIPE NCC, "RIPE Atlas: A Global Internet Measurement Network," in *Internet Protocol Journal (IPJ) '15*, Sep. 2015, http://ipj.dreamhosters.com/wp-content/uploads/2015/10/ipj18.3.pdf.

[14] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Comparing DNS Resolvers in the Wild," in *Internet Measurement Conference*, M. Allman, Ed. ACM, 2010, pp. 15–21. [Online]. Available: https://doi.org/10.1145/1879141.1879144

[15] J. S. Otto, M. A. Sánchez, J. P. Rula, and F. E. Bustamante, "Content Delivery and the Natural Evolution of DNS: Remote DNS Trends, Performance Issues and Alternative Solutions," in *Internet Measurement Conference*, 2012. [Online]. Available: https://doi.org/10.1145/2398776.2398831

[16] J. Damas, M. Graff, and P. Vixie, "Extension mechanisms for DNS (EDNS(0))," *RFC*, vol. 6891, pp. 1–16, 2013. [Online]. Available: https://doi.org/10.17487/RFC6891

[17] T. Callahan, M. Allman, and M. Rabinovich, "On Modern DNS Behavior and Properties," *Computer Communication Review*, vol. 43, no. 3, pp. 7–15, 2013. [Online]. Available: https://doi.org/10.1145/2500098.2500100

[18] G. Huston, "The Resolvers We Use," Nov. 2014, accessed 2021-Jan-21. [Online]. Available: https://labs.ripe.net/Members/gih/the-resolvers-we-use

[19] M. Allman, "Comments on DNS Robustness," in *Internet Measurement Conference*. ACM, 2018, pp. 84–90. [Online]. Available: https://doi.org/10.1145/3278532.3278541

[20] L. Zembruzki, A. S. Jacobs, G. S. Landtreter, L. Z. Granville, and G. C. M. Moura, "`dnstracker`: Measuring Centralization of DNS Infrastructure in the Wild," in *Advanced Information Networking and Applications*, ser. Advances in Intelligent Systems and Computing, vol. 1151. Springer, 2020, pp. 871–882. [Online]. Available: https://doi.org/10.1007/978-3-030-44041-1_76

[21] G. C. M. Moura, S. Castro, W. Hardaker, M. Wullink, and C. Hesselman, "Clouding up the Internet: how centralized is DNS traffic becoming?" in *Internet Measurement Conference*. ACM, 2020, pp. 42–49. [Online]. Available: https://doi.org/10.1145/3419394.3423625

[22] E. Targett, "Firefox Will Default to Cloudflare's Encrypted DNS-over-HTTPS Service," Sep. 2019, accessed 2021-Jan-21. [Online]. Available: https://www.cbronline.com/news/firefox-dns-over-https

[23] S. Deckelmann, "Mozilla Blog: Firefox continues push to bring DNS over HTTPS by default for US users," Feb. 2020, accessed 2021-Jan-21. [Online]. Available: https://blog.mozilla.org/blog/2020/02/25/firefox-continues-push-to-bring-dns-over-https-by-default-for-us-users/

[24] K. Baheux, "Chromium Blog: A safer and more private browsing experience with Secure DNS," May 2020, accessed 2021-Jan-21. [Online]. Available: https://blog.chromium.org/2020/05/a-safer-and-more-private-browsing-DoH.html

[25] D. Herrmann, C. Banse, and H. Federrath, "Behavior-based Tracking: Exploiting Characteristic Patterns in DNS Traffic," *Computers & Security*, vol. 39, pp. 17–33, 2013. [Online]. Available: https://doi.org/10.1016/j.cose.2013.03.012

[26] M. Kirchler, D. Herrmann, J. Lindemann, and M. Kloft, "Tracked Without a Trace: Linking Sessions of Users by Unsupervised Learning of Patterns in Their DNS Traffic," in *Workshop on Artificial Intelligence and Security*. ACM, 2016, pp. 23–34. [Online]. Available: https://doi.org/10.1145/2996758.2996770

[27] B. Greschbach, T. Pulls, L. M. Roberts, P. Winter, and N. Feamster, "The Effect of DNS on Tor's Anonymity," in *Network and Distributed System Security Symposium*. The Internet Society, 2017. [Online]. Available: https://www.ndss-symposium.org/ndss2017/ndss-2017-programme/e-effect-dns-tors-anonymity/

[28] M. Sun, G. Xu, J. Zhang, and D. W. Kim, "Tracking You through DNS Traffic: Linking User Sessions by Clustering with Dirichlet Mixture Model," in *Conference on Modelling, Analysis and Simulation of Wireless and Mobile Systems*. ACM, 2017, pp. 303–310. [Online]. Available: https://doi.org/10.1145/3127540.3127567

[29] A. Klein and B. Pinkas, "DNS Cache-Based User Tracking," in *Network and Distributed System Security Symposium*, 2019. [Online]. Available: https://www.ndss-symposium.org/ndss-paper/dns-cache-based-user-tracking/

[30] S. Bortzmeyer, "DNS privacy considerations," *RFC*, vol. 7626, pp. 1–17, 2015. [Online]. Available: https://doi.org/10.17487/RFC7626

[31] J. Arkko, B. Trammell, M. Nottingham, C. Huitema, M. Thomson, J. Tantsura, and N. ten Oever, "Considerations on Internet Consolidation and the Internet Architecture," Internet-Draft, Jul. 2019. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-arkko-iab-internet-consolidation-02

[32] Internet Society, "Internet Society Global Internet Report: Consolidation in the Internet Economy," 2019, accessed 2021-Jan-21. [Online]. Available: https://future.internetsociety.org/2019/

[33] Neustar, "Neustar Announces Acquisition of Verisign's Public DNS Service," Nov. 2020, accessed 2021-Jan-21. [Online]. Available: https://www.home.neustar/about-us/news-room/press-releases/2020/neustar-announces-acquisition-of-verisigns-public-dns-service

[34] V. Bajpai, S. J. Eravuchira, and J. Schönwälder, "Lessons Learned From Using the RIPE Atlas Platform for Measurement Research," *Computer Communication Review*, vol. 45, no. 3, pp. 35–42, 2015. [Online]. Available: https://doi.org/10.1145/2805789.2805796

[35] T. Holterbach, C. Pelsser, R. Bush, and L. Vanbever, "Quantifying Interference between Measurements on the RIPE Atlas Platform," in *Internet Measurement Conference*, 2015. [Online]. Available: https://doi.org/10.1145/2815675.2815710

[36] V. Bajpai, S. J. Eravuchira, J. Schönwälder, R. Kisteleki, and E. Aben, "Vantage Point Selection for IPv6 Measurements: Benefits and Limitations of RIPE Atlas Tags," in *Symposium on Integrated Network and Service Management*. IEEE, 2017, pp. 37–44. [Online]. Available: https://doi.org/10.23919/INM.2017.7987262

[37] Q. Scheitle, O. Hohlfeld, J. Gamba, J. Jelten, T. Zimmermann, S. D. Strowes, and N. Vallina-Rodriguez, "A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists," in *Internet Measurement Conference*, 2018, pp. 478–493. [Online]. Available: https://doi.org/10.1145/3278532.3278574

[38] E. P. Lewis, "The Role of Wildcards in the Domain Name System," *RFC*, vol. 4592, pp. 1–20, 2006. [Online]. Available: https://doi.org/10.17487/RFC4592

[39] B. Liu, C. Lu, H. Duan, Y. Liu, Z. Li, S. Hao, and M. Yang, "Who Is Answering My Queries: Understanding and Characterizing Interception of the DNS Resolution Path," in *USENIX Security Symposium*. USENIX Association, 2018, pp. 1113–1128. [Online]. Available: https://www.usenix.org/conference/usenixsecurity18/presentation/liu-baojun

[40] G. C. M. Moura, J. S. Heidemann, R. de Oliveira Schmidt, and W. Hardaker, "Cache Me If You Can: Effects of DNS Time-to-Live," in *Internet Measurement Conference*. ACM, 2019, pp. 101–115. [Online]. Available: https://doi.org/10.1145/3355369.3355568

[41] Y. Rekhter, B. G. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," *RFC*, vol. 1918, pp. 1–9, 1996. [Online]. Available: https://doi.org/10.17487/RFC1918

[42] R. de Oliveira Schmidt, J. S. Heidemann, and J. H. Kuipers, "Anycast Latency: How Many Sites Are Enough?" in *Passive and Active Measurement Conference*, vol. 10176. Springer, 2017, pp. 188–200. [Online]. Available: https://doi.org/10.1007/978-3-319-54328-4_14

[43] D. Atkins and R. Austein, "Threat Analysis of the Domain Name System (DNS)," *RFC*, vol. 3833, pp. 1–16, 2004. [Online]. Available: https://doi.org/10.17487/RFC3833

[44] T. Arnold, J. He, W. Jiang, M. Calder, Í. Cunha, V. Giotsas, and E. Katz-Bassett, "Cloud Provider Connectivity in the Flat Internet," in *Internet Measurement Conference*. ACM, 2020, pp. 230–246. [Online]. Available: https://doi.org/10.1145/3419394.3423613

[45] G. C. M. Moura, J. S. Heidemann, M. Müller, R. de Oliveira Schmidt, and M. Davids, "When the Dike Breaks: Dissecting DNS Defenses During DDoS," in *Internet Measurement Conference*. ACM, 2018, pp. 8–21. [Online]. Available: https://doi.org/10.1145/3278532.3278534