Can Encrypted DNS Be Fast?

Austin Hounsel[1(✉)], Paul Schmitt[1], Kevin Borgolte[2], and Nick Feamster[3]

[1] Princeton University, Princeton, NJ 08544, USA
{ahounsel,pschmitt}@cs.princeton.edu
[2] TU Delft, 2628 BX Delft, The Netherlands
k.borgolte@tudelft.nl
[3] University of Chicago, Chicago, IL 60637, USA
feamster@uchicago.edu

Abstract. In this paper, we study the performance of encrypted DNS protocols and conventional DNS from thousands of home networks in the United States, over one month in 2020. We perform these measurements from the homes of 2,693 participating panelists in the Federal Communications Commission's (FCC) Measuring Broadband America program. We found that clients do not have to trade DNS performance for privacy. For certain resolvers, DoT was able to perform *faster* than DNS in median response times, even as latency increased. We also found significant variation in DoH performance across recursive resolvers. Based on these results, we recommend that DNS clients (*e.g.*, web browsers) should periodically conduct simple latency and response time measurements to determine which protocol and resolver a client should use. No single DNS protocol nor resolver performed the best for all clients.

Keywords: DNS · Privacy · Security · Performance

1 Introduction

The Domain Name System (DNS) is responsible for translating human-readable domain names (*e.g.*, `nytimes.com`) to IP addresses. It is a critical part of the Internet's infrastructure that users must interact with before almost any communication can occur. For example, web browsers may require tens to hundreds of DNS requests to be issued before a web page can be loaded. As such, many design decisions for DNS have focused on minimizing the response times for requests. These decisions have in turn improved the performance of almost every application on the Internet.

In recent years, privacy has become a significant design consideration for the DNS. Research has shown that conventional DNS traffic can be passively observed by network eavesdroppers to infer which websites a user is visiting [2, 25]. This attack can be carried out by anyone that sits between a user and their recursive resolver. As a result, various protocols have been developed to send DNS queries over encrypted channels. Two prominent examples are DNS-over-TLS (DoT) and DNS-over-HTTPS (DoH) [8,10]. DoT establishes a TLS session

O. Hohlfeld et al. (Eds.): PAM 2021, LNCS 12671, pp. 444–459, 2021.
https://doi.org/10.1007/978-3-030-72582-2_26

over port 853 between a client and a recursive resolver. DoH also establishes a TLS session, but unlike DoT, all requests and responses are encoded in HTTP packets, and port 443 is used. In both cases, a client sends DNS queries to a recursive resolver over an encrypted transport protocol (TLS), which in turn relies on the Transmission Control Protocol (TCP). Encrypted DNS protocols prevent eavesdroppers from passively observing DNS traffic sent between users and their recursive resolvers. From a privacy perspective, DoT and DoH are attractive protocols, providing confidentiality guarantees that DNS lacked.

Past work has shown that typical DoT and DoH query response times are typically marginally slower than DNS [3,9,14]. However, these measurements were performed from university networks, proxy networks, and cloud data centers, rather than directly from homes. It is crucial to measure DNS performance from home networks in situ, as they may be differently connected than other networks. An early study on encrypted DNS performance was conducted by Mozilla at-scale with real browser users, but they did not study DoT, and they did not explore the effects of latency to resolvers, throughput, or Internet service provider (ISP) choice on performance [15]. Thus, the lack of controlled measurements prevents the networking community from fully understanding how encrypted DNS protocols perform for real users.

In this work, we provide a large-scale performance study of DNS, DoT, and DoH from thousands of home networks dispersed across the United States. We perform measurements from the homes of 2,693 participating panelists in the Federal Communications Commission's (FCC) Measuring Broadband America program from April 7th, 2020 through May 8th, 2020. We measure query response times and connection setup times using popular, open recursive resolvers, as well as resolvers provided by local networks. We also study the effects of latency to resolvers, throughput, and ISP choice on query response times.

2 Method

In this section, we describe the measurement platform we used to study DNS, DoT, and DoH performance and outline our analyses. We then describe the experiments we conduct and their limitations.

2.1 Measurement Platform

The FCC contracts with SamKnows [20] to implement the operational and logistical aspects of the Measuring Broadband America (MBA) program [6]. SamKnows is a company that specializes in developing custom software and hardware (also known as "Whiteboxes") to evaluate the performance of broadband access networks. Whiteboxes act as Ethernet bridges that connect directly to existing modems/routers, which enables us to control for poor Wi-Fi signals and cross-traffic. In accordance with MBA program objectives, SamKnows has deployed Whiteboxes to thousands of volunteers' homes across the United States.

We were granted access to the MBA platform through the FCC's MBA-Assisted Research Studies program (MARS) [5], which enables researchers (generally from the United States) to run measurements from the Whiteboxes. We utilize the platform to evaluate how DNS, DoT, and DoH perform from home networks.

We perform measurements from each Whitebox using SamKnows' DNS query tool. For each query, the tool reports a success/failure status (and failure reason, if applicable), the DNS resolution time excluding connection establishment (if the query was successful), and the resolved record [19]. For DoT and DoH, the tool separately reports the TCP connection setup time, the TLS session establishment time, and the DoH resolver lookup time. For this study, we only study queries for 'A' and 'AAAA' records. We note that queries for DNS and DoT are sent synchronously, *i.e.*, they must each receive a response before the next query can be sent. On the other hand, DoH queries are sent asynchronously, functionality that is enabled by the underlying HTTP protocol.

The query tool handles failures in several ways. First, if a response with an error code is returned from a recursive resolver (*e.g.*, NXDOMAIN or SERVFAIL), then the matching query is marked as a failure. Second, if the tool fails to establish a DoT or DoH connection, then all queries in the current batch (explained in Sect. 2.3) are marked as failures. Third, the query tool times out conventional DNS queries after three seconds, at which point it re-sends them. If three timeouts occur for a given query, the tool marks the query as a failure. Finally, the query tool marks DoT/DoH queries as failures if either five seconds have passed or if TCP hits the maximum number of re-transmissions allowed by the operating system's kernel (Linux 4.4.79). The Whiteboxes we measure use the default TCP settings configured by the kernel (*e.g.*, *net.ipv4.tcp_frto = 2*, *net.ipv4.tcp_retries1 = 3*, and *net.ipv4.tcp_retries2 = 15*).

In total, we collected measurements from 2,804 Whiteboxes, each of which use the latest generation of hardware and software (8.0) [21]. Our measurements were performed continuously from April 7th, 2020 through May 8th, 2020 in collaboration with SamKnows and the FCC. We filtered out certain Whiteboxes from our analysis in several ways. First, we filtered out 56 Whiteboxes that we did not have *any* network configuration information about (*e.g.*, ISP speed tier, ISP name, and access technology). Second, we filtered out 25 Whiteboxes that were connected by satellite. Third, we filtered out 30 Whiteboxes for which we did not know the access technology or ISP speed tier. This left us with 2,693 Whiteboxes to analyze, with 96% of queries marked as successful. The Whiteboxes were connected to 14 ISPs over cable, DSL, and fiber.

2.2 Analyses

We studied DNS, DoT, and DoH performance across several dimensions: connection setup times, query response times for each resolver and protocol, and query response times relative to latency to resolvers, throughput, and ISPs. Our analyses are driven by choices that DNS clients are able to make (*e.g.*, which protocol and resolver to use) and how these choices affect DNS performance.

Connection Setup Times. Before any query can be issued for DoT or DoH, the client must establish a TCP connection and a TLS session. As such, we measure the time to complete a 3-way TCP handshake and a TLS handshake. Additionally, for DoH, we measure the time to resolve the domain name of the resolver itself. The costs associated with connection establishment are amortized over many DoT or DoH queries as the connections are kept alive and used repeatedly once they are open. We study connection setup times in Sect. 3.1.

DNS Response Times. Query response times are crucial for determining the performance of various applications. Before any resource can be downloaded from a server, a DNS query often must be performed to learn the server's IP address (assuming a response is not cached). As such, we study query response times for each resolver and protocol in Sect. 3.2. We remove TCP and TLS connection establishment time from DoT and DoH query response times. The DNS query tool we use closes and re-establishes connections after ten queries (detailed in Sect. 2.3). As this behavior is unlikely to mimic that of stub resolvers and web browsers [7,16,17], we remove connection establishment times to avoid negatively biasing the performance of DoT and DoH.

DNS Response Times Relative to Latency and Throughput. Conventional DNS performance depends on latency, as the protocol is relatively lightweight; therefore, latency to the DNS resolver can have a significant effect on overall performance. Furthermore, encrypted DNS protocols may perform differently than conventional DNS in response to higher latency, as they are connection-oriented protocols. We study the effect of latency on query response times for each open resolver and protocol in Sect. 3.3. SamKnows also provides us with the subscribed downstream and upstream throughput for each Whitebox. We use this information to study the effect of downstream throughput on query response times in Sect. 3.3.

DNS Response Times Relative to ISP Choice. Lastly, SamKnows provides us with the ISP for each Whitebox. We study query response times for a selection of ISPs in Sect. 3.4.

2.3 Experiment Design

We describe below which recursive resolvers and domain names we perform measurements with and how we arrived at these choices.

DNS Resolvers. For each Whitebox, we perform measurements using three popular open recursive DNS resolvers (anonymized as X, Y, and Z, respectively[1]), as well as the recursive resolver automatically configured on each Whitebox (the "default" resolver). Typically, the default resolver is set by the ISP that

[1] We anonymize the resolvers as per our agreement with the FCC.

Table 1. Recursive resolver latency characteristics.

Resolver	Observations	Latency (ms)			
		Minimum	Median	Maximum	Std dev
X DNS and DoT	1,593,506	0.94	20.38	5,935.80	43.61
X DoH	1,567,337	0.14	22.75	8,929.88	43.25
Y DNS and DoT	1,596,964	2.00	20.90	9,701.82	46.79
Y DoH	1,552,595	0.14	20.50	10,516.31	40.68
Z DNS and DoT	1,579,605	2.35	31.41	516,844.73	414.26
Z DoH	1,533,380	0.14	33.00	9,537.42	41.11
Default DNS	2,009,086	0.13	0.85	8,602.39	22.93

the Whitebox is connected to. Resolvers X, Y, and Z all offer public name resolution for DNS, DoT, and DoH. However, the default resolver typically only supports DNS. As such, for the default resolver, we only perform measurements with conventional DNS. If a Whitebox has configured Resolver X, Y, or Z as its default resolver, then we leave its default resolver measurements out of our analysis.

In Table 1, we include the latency to each resolver across all Whiteboxes. We measure latency by running five ICMP ping tests for each resolver at the top of each hour and computing the average. We separate latency to DoH resolvers from latency to DNS and DoT resolvers because the domain names of DoH resolvers must be resolved in advance. As such, the IP addresses for the DoH resolvers are not always the same as DNS and DoT resolvers. We note that the latencies for the default resolvers are particularly low because these resolvers are often DNS forwarders configured on home routers. We exclude measurements with five failures or with an average latency of zero (0.7% of the total measurements).

We identified 41 Whiteboxes with median latencies to Resolvers X, Y, and Z DNS of up to 100 ms, despite median query response times of less than 1 ms. We consulted with SamKnows, and based on their experience, they believed this behavior could be attributed to DNS interception by middleboxes between Whiteboxes and recursive resolvers. For example, customer-premises equipment (CPE) can run DNS proxies (*e.g.*, dnsmasq) that can cache DNS responses to achieve such low query response times. Furthermore, previous reports from the United Kingdom indicate that ISPs can provide customer-premises equipment that is capable of passively observing and interfering with DNS queries [11]. We found that 29 of these 41 Whiteboxes are connected to the same ISP. We also identified two Whiteboxes with median latencies to X, Y, and Z DoH of less than 1 ms. Lastly, we identified one Whitebox with median latencies to X, Y, and Z DoT of up to 100 ms, despite median query response times of less than 1 ms. We analyze the data for these Whiteboxes for completeness.

Domain Names. Our goal was to collect DNS query response times for domain names found in websites that users are likely to visit. We first selected the top 100 websites in the Tranco top-list, which averages the rankings of websites in the

Alexa top-list over time [13]. For each website selected, we extracted the domain names of all included resources found on the page. We obtained this data from HTTP Archive Objects (or "HARs") that we collected from a previous study [9].

Importantly, we needed to ensure that the domain names were not sensitive in nature (*e.g.*, `pornhub.com`) so as to not trigger DNS-based parental controls. As such, after we created our initial list of domain names, we used the Webshrinker API to filter out domains associated with adult content, illegal content, gambling, and uncategorized content [24]. We then manually reviewed the resulting list. In total, our list included 1,711 unique domain names.[2]

Measurement Protocol. The steps we take to measure query response times from each Whitebox are as follows:

1. We randomize the input list of 1,711 domain names at the start of each hour.
2. We compute the latency to each resolver with a set of five ICMP ping tests.
3. We begin iterating over the randomized list by selecting a batch containing ten domain names.
4. We issue queries for all 10 domain names in the batch to each resolver/protocol combination. For DoT and DoH, we re-use the TLS connection for each query in the batch, and then close the connection. If a batch of queries has not completed within 30 s, we pause, check for cross-traffic, and retry if cross-traffic is present. If there is no cross traffic, we move to the next resolver/protocol combination.
5. We select the next batch of 10 domain names. If five minutes have passed, we stop for the hour. Otherwise, we return to step four.

Limitations. Due to bandwidth usage concerns and limited computational capabilities on the Whiteboxes, we do not collect web page load times while varying the underlying DNS protocol and resolver. Additionally, while we conducted our measurements, the COVID-19 pandemic caused many people to work from home. We did not want to perturb other measurements being run with the Measuring Broadband America platform or introduce excessive strain on the volunteers' home networks. Due to these factors, we focus on DNS response times.

3 Results

This section presents the results of our measurements. We organize our results around the following questions: (1) How much connection overhead does encrypted DNS incur, in terms of resolver lookup (in the case of DoH), TCP connect time, and TLS setup time; (2) How does encrypted DNS perform versus conventional DNS?; (3) How does network performance affect encrypted DNS

[2] Our list of domain names that we measured is available at https://github.com/noise-lab/dns-mba-public.git.

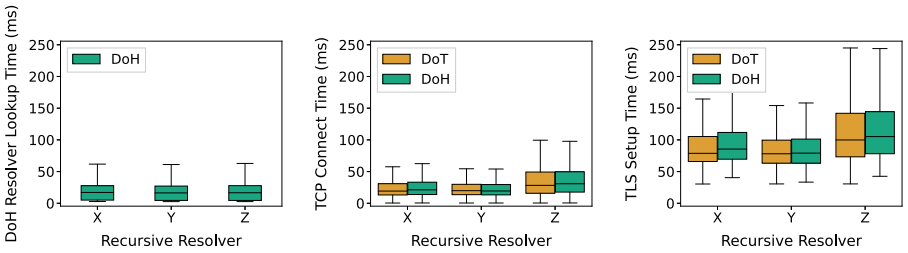

(a) DoH Resolver Lookup (b) TCP Connect Time (c) TLS Setup Time

Fig. 1. Connection setup times for DoT and DoH.

performance?; and (4) How does encrypted DNS resolver performance depend on broadband access ISP? Our results show that in the case of certain resolvers—to our surprise—DoT had *lower* median response times than conventional DNS, even as latency to the resolver increased. We also found significant variation in DoH performance across resolvers.

3.1 How Much Connection Overhead Does Encrypted DNS Incur?

We first study the overhead incurred by encrypted DNS protocols, due to their requirements for TCP connection setup and TLS handshakes. Before any batch of DoT queries can be issued with the SamKnows query tool, a TCP connection and TLS session must be established with a recursive resolver. In the case of DoH, the resolver's domain name is also resolved (*e.g.*, `resolverX.com`). In Fig. 1, we show timings for different aspects of connection establishment for DoT and DoH. The results show that lookup times were similar for all three resolvers (Fig. 1(a)). This result is expected because the same default, conventional DNS resolver is used to look up the DoH resolvers' domain names; the largest median DoH resolver lookup time was X with 17.1 ms. Depending on the DNS time to live (TTL) of the DoH resolver lookup, resolution of the DoH resolver may occur frequently or infrequently.

Next, we study the TCP connection establishment time for DoT and DoH for each of the three recursive resolvers (Fig. 1(b)). For each of the three individual resolvers, TCP establishment time for DoT and DoH are similar. Resolvers X and Y are similar; Z experienced longer TCP connection times. The largest median TCP connection establishment time across all resolvers and protocols (Resolver Z DoH) was 30.8 ms.

Because DoT and DoH rely on TLS for encryption, a TLS session must be established before use. Figure 1(c) shows the TLS establishment time for the three open resolvers. Again, Resolver Z experienced higher TLS setup times compared to X and Y. Furthermore, DoT and DoH performed similarly for each resolver. The largest median TLS connection establishment time across all recursive resolvers and protocols (Resolver Z DoH) was 105.2 ms. As with resolver lookup overhead, the cost of establishing a TCP and TLS connection to

the recursive resolver for a system would ideally occur infrequently, and should be amortized over many queries by keeping the connection alive and reusing it for multiple DNS queries.

Connection-oriented, secure DNS protocols will incur additional latency, but these costs can be (and are) typically amortized by caching the DNS name of the DoH resolver, as well as multiplexing many DNS queries over a single TLS session to a DoH resolver. Many browser implementations of DoH implement these practices. For example, Firefox establishes a DoH connection when the browser launches, and it leaves the connection open [16,17]. Thus, the overhead for DoH connection establishment in Firefox is amortized over time.

In the remainder of this paper we do not include connection establishment overhead when studying DNS query response times. We omit connection establishment time for the rest of our analysis because the DNS query tool closes and re-opens connections for each batch of queries. Thus, inclusion of TCP and TLS connection overheads may negatively skew query response times.

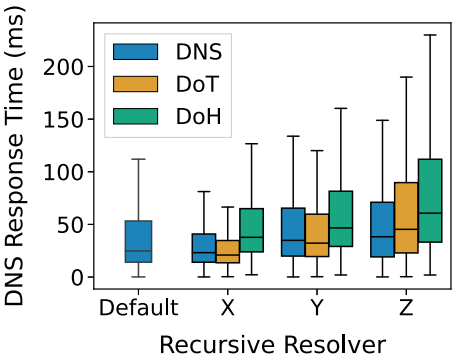

Fig. 2. Aggregate query response times.

3.2 How Does Encrypted DNS Perform Compared with Conventional DNS?

We next compare query response times across each protocol and recursive resolver. Figure 2 shows box plots for DNS response times across all Whiteboxes for each resolver and protocol. "Default" refers to the resolver that is configured by default on each Whitebox (which is typically the DNS resolver operated by the Whitebox's upstream ISP).

DNS Performance Varies Across Resolvers. First of all, conventional DNS performance varies across recursive resolvers. For the default resolvers configured on Whiteboxes, the median query response time using conventional DNS is 24.8 ms. For Resolvers X, Y, and Z, the median query response times using DNS are

23.2 ms, 34.8 ms, and 38.3 ms, respectively. Although X performs better than the default resolvers, Y and Z perform at least 10 ms slower. This variability could be attributed to differences in deployments between open resolvers.

DoT Performance Nearly Matches Conventional DNS. Interestingly DoT lookup times are close to those of conventional DNS. For Resolvers X, Y, and Z, the median query response times for DoT are 20.9 ms, 32.2 ms, and 45.3 ms, respectively. Interestingly, for X and Y, we find that DoT performs 2.3 ms and 2.6 ms *faster* than conventional DNS, respectively. For both of these resolvers, the best median DNS query performance could be attained using DoT. Z's median response time was 7 ms slower. The performance improvement of DoT over conventional DNS in some cases is interesting because conventional wisdom suggests that the connection overhead of TCP and TLS would be prohibitive. On the other hand, various factors, including transport-layer optimizations in TCP, as well as differences in infrastructure deployments, could explain these discrepancies. It may also be the case that DoT resolvers have lower query loads than conventional DNS resolvers, enabling comparable (or sometimes faster) response times. Investigating the causes of these discrepancies is an avenue for future work.

DoH Response Times were Higher Than Those for DNS and DoT. DoH experienced higher response times than conventional DNS or DoT, although this difference in performance varies significantly across DoH resolvers. For Resolvers X, Y, and Z, the median query response times for DoH are 37.7 ms, 46.6 ms, and 60.7 ms, respectively. Resolver Z exhibited the biggest increase in response latency between DoH and DNS (22.4 ms). Resolver Y showed the smallest difference in performance between DoH and DNS (11.8 ms). Median DoH response times between resolvers can differ greatly, with X DoH performing 23 ms faster than Z DoH. The performance cost of DoH may be due to the overhead of HTTPS, as well as the fact that DoH implementations are still relatively nascent, and thus may not be optimized. For example, an experimental DoH recursive resolver implementation by Facebook engineers terminates DoH connections to a reverse web proxy before forwarding the query to a DNS resolver [4].

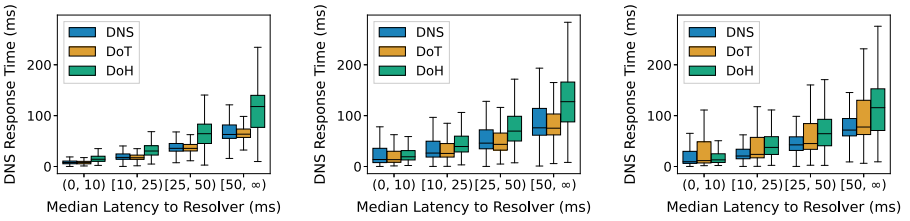

(a) Resolver X (b) Resolver Y (c) Resolver Z

Fig. 3. DNS response times based on median latency to resolvers.

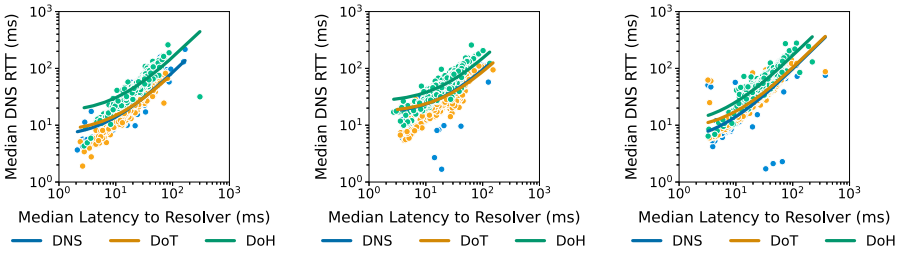

(a) Resolver X (b) Resolver Y (c) Resolver Z

Fig. 4. Ridge regression models comparing median latency to resolvers to median DNS response times (alpha = 1).

Table 2. Coefficients, intercepts, and errors for ridge regression models.

Resolver	Coefficient	Intercept	Mean absolute error	Mean squared error
X DNS	0.79	6.01	3.70	62.06
X DoT	0.74	7.48	4.23	33.89
X DoH	1.41	16.39	11.82	551.74
Y DNS	0.79	15.57	8.35	109.25
Y DoT	0.71	16.67	9.20	126.43
Y DoH	1.26	25.17	12.36	289.20
Z DNS	0.93	4.82	4.46	221.03
Z DoT	0.95	8.07	5.58	221.91
Z DoH	1.59	9.75	14.29	482.44

3.3 How Does Network Performance Affect Encrypted DNS Performance?

We next study how network latency and throughput characteristics affect the performance of encrypted DNS.

DoT Can Meet or Beat Conventional DNS Despite High Latencies to Resolvers, Offering Privacy Benefits for no Performance Cost. Figure 3 shows that DoT can perform better than DNS as latency increases for Resolvers X and Y; in the case of Resolver Z, DoT nearly matches the performance of conventional DNS. We observe similar behavior with the linear ridge regression models shown in Fig. 4. As discussed in Sect. 3.2, these results could be explained by transport-layer optimizations in TCP, differences in infrastructure deployments, and lower query loads on DoT resolvers compared to conventional DNS resolvers.

DoH Performs Worse Than Conventional DNS and DoT as Latencies To Resolvers Increase. Figure 3 shows that DoH performs substantially worse when latency between the client and recursive resolver is high; Fig. 4 shows a similar result with a ridge regression model. As discussed in Sect. 3.2, this result could be explained by either HTTPS overhead, nascent DoH implementations and deployments, or both.

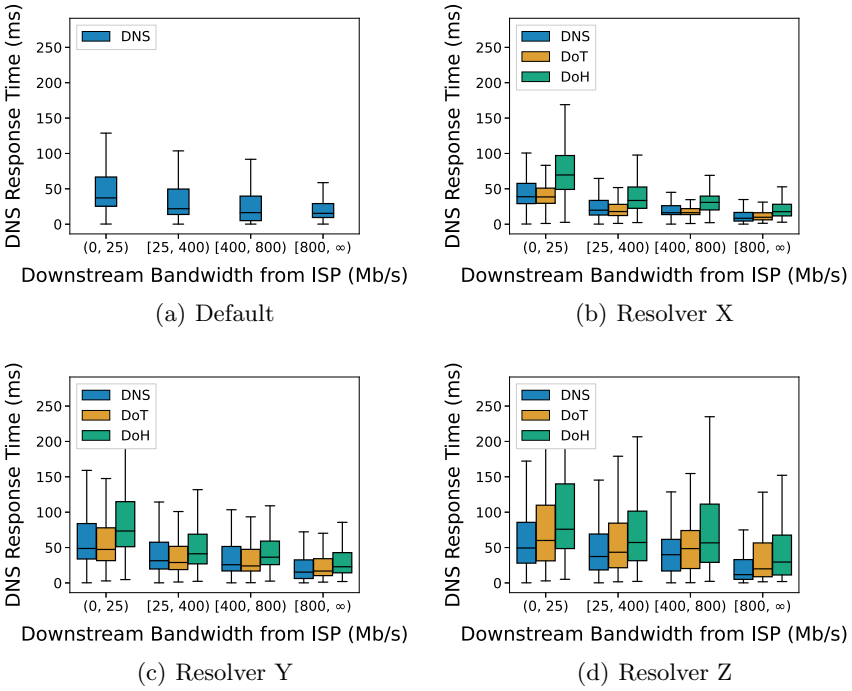

(a) Default (b) Resolver X

(c) Resolver Y (d) Resolver Z

Fig. 5. Query response times based on downstream access ISP throughput.

Subscribed Throughput Affects DNS Performance. Figure 5 shows DNS response times across each of the open resolvers as well as the default resolver. We bin the downstream throughput into four groups using clustering based on kernel density estimation. The performance for all protocols tends to improve as throughput increases, with DoH experiencing the most relative improvement. For example, for users with throughput that is less than 25 Mbps, the median query response times for Resolver Y DoH and Y DNS are 73.4 ms and 48.7 ms, respectively. As throughput increases from 25 Mbs to 400 Mbps, the median query response times for Y DoH and Y DNS are 41.2 ms and 31.4 ms, respectively. DoT performs similarly to conventional DNS regardless of downstream throughput. Across all groups, the absolute performance difference between Resolver X DoT and X DNS by 0.2 ms, 1.9 ms, 0.1 ms, and 1.4 ms, respectively. For Resolver Y, DoT again

performs faster than DNS in median query response times when throughput is less than 800 Mbps. For the three lower throughput groups, Y DoT performs faster than Y DNS by 1.4 ms, 2.5 ms, and 1.7 ms, respectively.

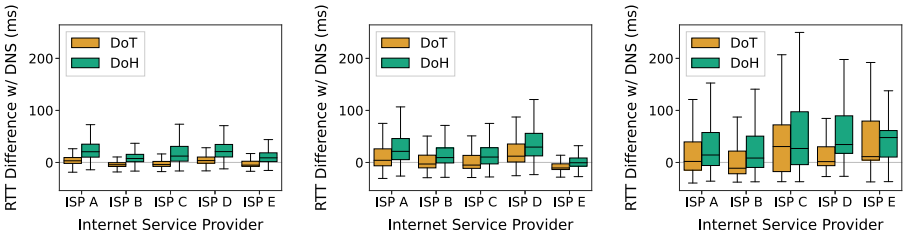

(a) Resolver X (b) Resolver Y (c) Resolver Z

Fig. 6. Per-ISP query response times.

3.4 Does Encrypted DNS Resolver Performance Vary Across ISPs?

Figure 6 shows how encrypted DNS response times vary across different resolvers and ISPs. In short, the choice of resolver matters, and the "best" encrypted DNS resolver also may depend on the user's ISP. For instance, while ISP C is comparable to the other ISPs for queries sent to Resolver X, ISP C has significantly lower query response times to Resolver Y, and is one of the poorest performing ISPs on Resolver Z. The difference in median query response times between Resolver X DoH and X DNS was 20.9 ms for Whiteboxes on ISP D, and 8.9 ms for Whiteboxes on ISP E; for Z DoH, the difference in median times was 34.5 ms for Whiteboxes on ISP D, and 47.9 ms for Whiteboxes on ISP E.

Resolver performance can also differ across ISPs. For ISP B, the median query response time for Z DoT is 11.1 ms faster than Z DNS. However, for ISP C, Z DoT is significantly slower than DNS, with a difference in median query response times of 30.6 ms. We attribute this difference in performance to higher latency to Resolver Z via ISP C. The median latency to Z DNS and DoT across Whiteboxes on ISP C was 50 ms, compared to 18.5 ms on ISP B.

4 Related Work

Researchers have compared the performance of DNS, DoT, and DoH in various ways. Zhu et al. proposed DoT to encrypt DNS traffic between clients and recursive resolvers [25]. They modeled its performance and found that DoT's overhead can be largely eliminated with connection re-use. Böttger et al. measured the effect of DoT and DoH on query response times and page load times from a university network [3]. They find that DNS generally outperforms DoT in response times, and DoT outperforms DoH. Hounsel et al. also measure response times and page load times for DNS, DoT, and DoH using Amazon EC2 instances [9].

They find that despite higher response times, page load times for DoT and DoH can be *faster* than DNS on lossy networks. Lu et al. utilized residential TCP SOCKS networks to measure response times from 166 countries and found that, in the median case with connection re-use, DoT and DoH were slower than conventional DNS over TCP by 9 ms and 6 ms, respectively [14].

Researchers have also studied in depth how DNS influences application performance. Sundaresan et al. used an early MBA deployment of 4,200 home gateways to identify performance bottlenecks for residential broadband networks [22]. This study found that page load times for users in home networks are significantly influenced by slow DNS response times. Wang et al. introduced WProf, a profiling system that analyzes various factors that contribute to page load times [23]. They found that queries for uncached domain names at recursive resolvers can account for up to 13% of the critical path delay for page loads. Otto et al. found that CDN performance was significantly affected by clients choosing recursive resolvers that are far away from CDN caches [18]. As a result of these findings. Otto et al. proposed *namehelp*, a DNS proxy that sends queries for CDN-hosted content to directly to authoritative servers. Allman studied conventional DNS performance from 100 residences in a neighborhood and found that only 3.6% of connections were blocked on DNS with lookup times greater than either 20 ms or 1% of the application's total transaction time [1].

Past work studied the performance impact of "last mile" connections to home networks in various ways. Kreibich et al. proposed Netalyzr as a Java applet that users run from devices in their home networks to test debug their Internet connectivity. Netalyzr probes test servers outside of the home network to measure latency, IPv6 support, DNS manipulation, and more. Their system was run from over 99,000 public IP addresses, which enabled them to study network connectivity at scale [12]. Dischinger et al. measured bandwidth, latency, and packet loss from 1,894 hosts and 11 major commercial cable and DSL providers in North America and Europe. This work found that the "last mile" connection between an ISP and a home network is often a performance bottleneck, which they could not have captured by performing measurements outside of the home network. However, their measurements were performed from hosts located within homes, rather than the home gateway. This introduces confounding factors between hosts and the home gateway, such as poor Wi-Fi performance.

5 Conclusion

In this paper, we studied the performance of encrypted DNS protocols and DNS from 2,693 Whiteboxes in the United States, between April 7th, 2020 and May 8th, 2020. We found that clients do not have to trade DNS performance for privacy. For certain resolvers, DoT was able to perform *faster* than DNS in median response times, even as latency increased. We also found significant variation in DoH performance across recursive resolvers. Based on these results, we recommend that DNS clients (*e.g.*, web browsers) measure latency to resolvers and DNS response times determine which protocol and resolver a client should use. No single DNS protocol nor resolver performed the best for all clients.

There were some limitations to our work that point to future research. First, due to bandwidth restrictions, we were unable to perform page loads from Whiteboxes. Future work could utilize platforms of similar scale to SamKnows to measure page loads, such as browser telemetry systems. Second, future work should perform measurements from mobile devices. DoT was implemented in Android 10, but to our knowledge, its performance has not been studied "in the wild." Finally, future work could study how encrypted DNS protocols perform from networks that are far away from popular resolvers. This is particularly important for browser vendors that seek to deploy DoH outside of the United States.

Acknowledgements. We thank the Federal Communications Commission's Measuring Broadband America (MBA) program and the associated MBA-Assisted Research (MARS) Program for assistance in conducting this study, Jason Livingood and Al Morton for initial study design suggestions, the MBA collaborative for experiment input, and Sam Crawford from SamKnows with assistance in measurement implementation and deployment. This research was funded in part by National Science Foundation Award CNS-1704077 and a Comcast Innovation Fund.

References

1. Allman, M.: Putting DNS in context. In: Chritin, N., Pelechrinis, K., Sekar, V. (eds.) Proceedings of the 2020 Internet Measurement Conference (IMC). Association for Computing Machinery (ACM) (2020)
2. Bortzmeyer, S.: DNS Privacy Considerations. RFC 7626, RFC Editor (2015). http://www.ietf.org/rfc/rfc7626.txt. (Informational)
3. Böttger, T., et al.: An empirical study of the cost of DNS-over-https. In: Sperotto, A., van Rijswijk-Deij, R., Hesselman, C. (eds.) Proceedings of the 2019 Internet Measurement Conference, Amsterdam, Netherlands, pp. 15–21. Association for Computing Machinery (ACM) (2019). https://doi.org/10.1145/3355369.3355575. https://dl.acm.org/doi/pdf/10.1145/3355369.3355575
4. Facebook Experimental: Doh proxy (2020). https://facebookexperimental.github.io/doh-proxy/
5. Federal Communications Commission: MBA Assisted Research Studies (2020). https://www.fcc.gov/general/mba-assisted-research-studies
6. Federal Communications Commission: Measuring Broadband America (2020). https://www.fcc.gov/general/measuring-broadband-america
7. getdns Team: getdns/stubby (2019). https://github.com/getdnsapi/stubby
8. Hoffman, P., McManus, P.: DNS Queries over HTTPS (DoH). RFC 8484, RFC Editor (2018). http://www.ietf.org/rfc/rfc8484.txt. (Proposed Standard)
9. Hounsel, A., Borgolte, K., Schmitt, P., Holland, J., Feamster, N.: Comparing the effects of DNS, dot, and DOH on web performance. In: Huang, Y., King, I., Liu, T.Y., van Steen, M. (eds.) Proceedings of the 28th The Web Conference (WWW), Taipei, Taiwan, pp. 562–572. Association for Computing Machinery (ACM) (2020). https://doi.org/10.1145/3366423.3380139. https://dl.acm.org/doi/pdf/10.1145/3366423.3380139
10. Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessel, D., Hoffman, P.: Specification for DNS over Transport Layer Security (TLS). RFC 7858, RFC Editor (2016). http://www.ietf.org/rfc/rfc7858.txt. (Proposed Standard)

11. Jackson, M.: Firmware update for UK sky broadband ISP routers botches DNS update (2019). https://www.ispreview.co.uk/index.php/2019/04/firmware-update-for-uk-sky-broadband-isp-routers-botches-dns.html
12. Kreibich, C., Weaver, N., Nechaev, B., Paxson, V.: Netalyzr: illuminating the edge network. In: Allman, M. (ed.) Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC), Melbourne, Australia, pp. 246–259. Association for Computing Machinery (ACM) (2010). https://doi.org/10.1145/1879141.1879173. https://dl.acm.org/doi/pdf/10.1145/1879141.1879173
13. L. Pochat, V., V. Goethem, T., Tajalizadehkhoob, S., Korczyński, M., Joosen, W.: Tranco: a research-oriented top sites ranking hardened against manipulation. In: Oprea, A., Xu, D. (eds.) Proceedings of the 26th Network and Distributed System Security Symposium (NDSS), San Diego, CA, USA, pp. 1–15. Internet Society (ISOC) (2019). https://doi.org/10.14722/ndss.2019.23386. https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_01B-3_LePochat_paper.pdf
14. Lu, C., et al.: An end-to-end, large-scale measurement of DNS-over-encryption: how far have we come? In: Sperotto, A., van Rijswijk-Deij, R., Hesselman, C. (eds.) Proceedings of the 2019 Internet Measurement Conference, Amsterdam, Netherlands, pp. 22–35. Association for Computing Machinery (ACM) (2019). https://doi.org/10.1145/3355369.3355580. https://dl.acm.org/doi/pdf/10.1145/3355369.3355580
15. McManus, P.: Firefox Nightly Secure DNS Experimental Results (2018). https://blog.nightly.mozilla.org/2018/08/28/firefox-nightly-secure-dns-experimental-results/
16. Mozilla: All.js (2020). https://searchfox.org/mozilla-central/source/modules/libpref/init/all.js#1425
17. Mozilla: TRRServiceChannel.cpp (2020). https://searchfox.org/mozilla-central/source/netwerk/protocol/http/TRRServiceChannel.cpp#512
18. Otto, J.S., Sánchez, M.A., Rula, J.P., Bustamante, F.E.: Content delivery and the natural evolution of DNS: remote DNS trends, performance issues and alternative solutions. In: Mahajan, R., Snoeren, A. (eds.) Proceedings of the 2012 Internet Measurement Conference (IMC), Boston, MA, USA, pp. 523–536. Association for Computing Machinery (ACM) (2012). https://doi.org/10.1145/2398776.2398831. https://dl.acm.org/doi/pdf/10.1145/2398776.2398831
19. SamKnows: DNS resolution (2020). https://samknows.com/technology/tests/dns-resolution
20. SamKnows: SamKnows (2020). https://www.samknows.com/
21. SamKnows: SamKnows Whitebox (2020). https://samknows.com/technology/agents/samknows-whitebox#specifications
22. Sundaresan, S., Feamster, N., Teixeira, R., Magharei, N.: Measuring and mitigating web performance bottlenecks in broadband access networks. In: Gummadi, K., Partidge, C. (eds.) Proceedings of the 2013 Internet Measurement Conference (IMC), Barcelona, Spain, pp. 213–226. Association for Computing Machinery (ACM) (2013). https://doi.org/10.1145/2504730.2504741. https://dl.acm.org/doi/10.1145/2504730.2504741
23. Wang, X.S., Balasubramanian, A., Krishnamurthy, A., Wetherall, D.: Demystifying page load performance with WProf. In: Feamster, N., Mogul, J. (eds.) Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI), Lombard, IL, USA, pp. 473–487. USENIX Association (2013). https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/wang_xiao

24. Webshrinker: APIs - Webshrinker (2020). https://www.webshrinker.com/apis/
25. Zhu, L., Hu, Z., Heidemann, J., Wessels, D., Mankin, A., Somaiya, N.: Connection-oriented DNS to improve privacy and security. In: Shmatikov, V., Bauer, L. (eds.) Proceedings of the 36th IEEE Symposium on Security & Privacy (S&P), San Jose, CA, USA, pp. 171–186. Institute of Electrical and Electronics Engineers (IEEE) (2015). https://doi.org/10.1109/sp.2015.18. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7163025