

Measuring DNS-over-HTTPS Performance Around the World

Rishabh Chhabra
University of Illinois at
Urbana-Champaign
IL, USA
chhabra4@illinois.edu

Paul Murley
University of Illinois at
Urbana-Champaign
IL, USA
pmurley2@illinois.edu

Deepak Kumar
Stanford University
CA, USA
kumarde@cs.stanford.edu

Michael Bailey
University of Illinois at
Urbana-Champaign
IL, USA
mdbailey@illinois.edu

Gang Wang
University of Illinois at
Urbana-Champaign
IL, USA
gangw@illinois.edu

ABSTRACT

In recent years, DNS-over-HTTPS (DoH) has gained significant traction as a privacy-preserving alternative to unencrypted DNS. While several studies have measured DoH performance relative to traditional DNS and other encrypted DNS schemes, they are often incomplete, either conducting measurements from single countries or are unable to compare encrypted DNS to default client behavior. To expand on existing research, we use the BrightData proxy network to gather a dataset consisting of 22,052 unique clients across 224 countries and territories. Our data shows that the performance impact of a switch to DoH is mixed, with a median slowdown of 65ms per query across a 10-query connection, but with 28% of clients receiving a speedup over that same interval. We compare four public DoH providers, noting that Cloudflare excels in both DoH resolution time (265ms) and global points-of-presence (146). Furthermore, we analyze geographic differences between DoH and Do53 resolution times, and provide analysis on possible causes, finding that clients from countries with low Internet infrastructure investment are almost twice as likely to experience a slowdown when switching to DoH as those with high Internet infrastructure investment. We conclude with possible improvements to the DoH ecosystem. We hope that our findings can help to inform continuing DoH deployments.

CCS CONCEPTS

• **Networks** → **Naming and addressing**; **Network measurement**.

ACM Reference Format:

Rishabh Chhabra, Paul Murley, Deepak Kumar, Michael Bailey, and Gang Wang. 2021. Measuring DNS-over-HTTPS Performance Around the World. In *Internet Measurement Conference (IMC '21)*, November 2–4, 2021, Virtual Event. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3487552.3487849>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '21, November 2–4, 2021, Virtual Event

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9129-0/21/11...\$15.00

<https://doi.org/10.1145/3487552.3487849>

1 INTRODUCTION

Over the past few years, several industry actors have advocated for a transition to DNS-over-HTTPS (DoH) as a privacy-preserving alternative to traditional, unencrypted, UDP-based DNS (Do53). Mozilla Firefox already defaults to DoH in Firefox for clients in the United States [36], Google has announced a gradual rollout of DoH by default in Google Chrome [6], Microsoft plans to build DoH into both the Edge browser and Windows operating system itself [34], and Apple has built encrypted DNS into their platforms for developers to integrate [1].

In response to these commitments to deploy DoH, prior work has investigated encrypted DNS performance in many ways. Hounsell et al. studied how encrypted DNS affects web browsing using five EC2 nodes [21] and the direct performance impact faced by US clients [22]. Lu et al. studied the reachability and performance of DoH servers to residential nodes around the world [29], though they approximated Do53 behavior using TCP and only with selected resolvers (instead of clients' default resolvers). Inspired by these efforts, our paper focuses on capturing direct comparisons between DoH performance and the *default* client DNS behavior to assess the performance impact of the transition to DoH. In addition, we seek to understand whether (and how) the transition to DoH would *unequally* affect clients in different countries and territories. Understanding these differences is key to understanding how to make equitable rollout decisions for DoH worldwide.

In this paper, we leverage BrightData (formerly Luminati) [4], a residential HTTPS proxy network, to conduct DoH and Do53 performance measurements from 22,052 clients located in 224 countries and territories. We develop careful heuristics for measuring DoH and Do53 performance through the BrightData network by instructing BrightData clients to resolve fresh domain names under our control (i.e., cache miss performance). This allows us to explore the performance lower-bound for both Do53 and DoH. We demonstrate through ground truth validation experiments that our heuristics almost exactly approximate DoH and Do53 performance for BrightData clients, with errors of up to 10ms for DoH and 2ms for Do53. Our results serve as the closest exact measurements of DoH and Do53 performance for residential clients around the world.

We compare DoH measurements drawn from four public resolution services (Cloudflare, Google, NextDNS, Quad9) and default resolution behavior on end-clients. We find that clients globally

take a median 415ms to resolve an initial DoH query, compared to a median 234ms for a single Do53 query. 19.1% of DoH clients enjoy a speedup in performance even on the first request (despite the TLS handshake), aligned with prior studies suggesting that DoH may outperform Do53 in select cases [20]. For example, clients in Indonesia see their median resolution time drop by 179ms upon switching from Do53 to DoH. Most clients and countries, however, do not enjoy such a speedup. Even after accounting for time spent on the initial TLS handshake, clients in Sudan, for example, experience a 264ms median increase in resolution time across the four DoH providers we study. We also examine differences between DoH providers, and finding that Cloudflare has 36% more points-of-presence (PoP) and resolves queries 21% faster than the next closest DoH resolution service. In addition, we approximate geographic distances between clients and their resolvers, finding that DoH providers often fail to select the closest PoP for each client, sometimes by huge margins. Quad9, for example, only assigns 21% of clients to the closest available PoP, according to our dataset.

To better explain differences in DoH performance around the world, we model DoH and Do53 performance as the outcome of several explanatory variables focused on Internet infrastructure investment, economic development, and DoH infrastructure properties (e.g., resolver choice, PoP placement). We find countries with low economic development and low Internet infrastructure investment are more likely to experience significant DoH slowdowns compared to Do53. For example, clients from countries with nationwide bandwidth <25Mbps experience a median slowdown of 350ms when transitioning to DoH from Do53, compared to just 112ms for clients in countries with faster Internet speeds. We observe these trends are still significant even when considering multiple requests using a single TLS session for DoH queries, highlighting that while reused connections may dampen the performance cost slightly, they still disproportionately impact countries with fewer economic capacity.

We conclude with a discussion of the implications of our measurements on DoH rollout globally and how to support and enable future research in this space. To this end, we provide our dataset¹ in the hopes that it may aid in further research. We hope our results will add context to the discussion surrounding DoH deployment and can inform relevant parties on DoH deployment strategies.

In summary, our contributions are as follows:

- (1) We conduct measurements of DoH and Do53 performance globally at 22,052 residential clients from 224 countries, and are able to directly evaluate the performance cost of switching to DoH from default client DNS behavior. Our dataset will be released at publication time.
- (2) We show that DoH performance varies for clients around the world, and that while most clients would experience only a moderate slowdown, 10% of the clients in our dataset see their resolution times triple as a result of switching from their default resolver to DoH. We find that 8.8% of the countries benefit from a switch to DoH from Do53.
- (3) We find that a significant number of clients are not being routed to the public resolver PoP nearest to them. For example, 26% of Cloudflare clients could be switched to a PoP at least 1,000 miles closer.
- (4) We model DoH and Do53 performance as an outcome of several explanatory variables, and find that countries with lower Internet infrastructure investment will experience disproportionate slowdowns in a unilateral switch to DoH from Do53.

2 BACKGROUND AND GOALS

In this section, we introduce relevant encrypted DNS work and describe our research questions.

DNS-over-Encryption. The Domain Name System (DNS) allows clients to look up a human-readable domain name to obtain its IP address [35]. The commonly used DNS protocol uses port 53 (referenced as “Do53”) and supports *unencrypted* queries over UDP and TCP. The fact the DNS remains widely unencrypted raises security and privacy concerns which have been well-studied, such as connection eavesdropping or tampering with DNS traffic [47]. To secure DNS, the Internet Engineering Task Force (IETF) has proposed and developed five major protocols: DNS-over-TLS (DoT), DNS-over-HTTPS (DoH), DNS-over-QUIC, DNSCrypt, and DNSSEC [47].

DoH vs. DoT. Among existing DNS-over-Encryption solutions, DoT and DoH have gained the widest adoption in practice [29]. Both protocols send DNS traffic over a TLS connection, with DoH sending queries in an HTTP GET request. Recent reports show that DoH has gained more traction than DoT [8, 25], in part because DoH causes fewer problems with port-oriented firewalls since it uses port 443 instead of alternate ports (DoT uses port 843 by default) [8]. This also makes DoH more robust to censorship [21], as a censor is unlikely to block port 443 universally. As DoH seems to be the most widespread encrypted DNS standard in use today, we focus primarily on it in this research.

Why Not Existing Methods? Existing measurement methodologies are insufficient to answer our research questions. The first type of methodology requires direct control over each vantage point. As a result, the number of vantage points is highly limited (e.g., a single machine [8] or 5 Amazon EC2 nodes [21]). A recent work obtained access to 2.6K volunteer nodes from the FCC Measuring Broadband America program, but all of these vantage points were in the United States [22]. Doan et al. measured the performance of DoT using 3.2K RIPE Atlas probes located in residential networks [16]. However, their measurement only covered DoT, and the methodology cannot be used for DoH measurement due to the API restrictions of RIPE Atlas.

Another technique in the literature made use of SOCKS proxy networks [29] to run DoH and Do53 measurements through a large number of vantage points. However, because of the lack of control over the vantage points and involvements of proxy servers (middleboxes), this method cannot obtain the *absolute query latency* for DoH or Do53 [29]. Instead, they can only obtain the differential. In addition, their measurement technique only supported DNS-over-TCP (instead of the more common DNS-over-UDP), and did not

¹Dataset: <https://github.com/rishabhch/imec21-measuring-doh-performance>.




Figure 1: Our Experimental Setup—The measurement client, web server, and authoritative name server are under our control. We use proxy service BrightData to reach a large number of clients (exit nodes) located in different countries. Each client sends queries over DNS-over-HTTPS (DoH) and conventional DNS (Do53), and we measure the timing and results.

allow them to measure the performance of default resolvers of their clients.

Different from existing works, our methodology aims to (1) cover a large number of residential vantage points from many countries, (2) study the behavior of default client resolvers *as configured* relative to DoH, and (3) obtain absolute query latency from both our DoH and Do53 measurements. While we do not expect DoH to outperform Do53 in terms of latency, we want to study the performance impact of a transition from Do53 to DoH as it may disproportionately impact some user populations in certain world regions compared to others.

The following related questions guide our research:

- (1) What is the performance impact on real-world clients of a switch from traditional Do53 to DoH?
- (2) How does this impact differ across countries and geographical regions?
- (3) What external factors or variables explain the performance asymmetries we observe?
- (4) How do public DoH services differ in their architecture, and how do these differences affect end clients?

3 METHODOLOGY

In this section, we describe our measurement methodology to achieve the goals described above.

3.1 Methodology Overview

To conduct DoH and Do53 measurements, we utilize a proxy service called BrightData (formerly known as Luminati) [4] to solicit measurement vantage points from a large number of countries. BrightData is a paid HTTPS proxy service that routes traffic globally via *exit nodes* that have HolaVPN installed [30, 46]. HolaVPN is a community powered VPN that gives users free VPN access in exchange for the users’ machines becoming part of the Hola network. A key advantage of BrightData is that it allows us to perform measurements via machines located in *residential networks*. For an extended discussion of the ethics of using proxy networks like BrightData as measurement platforms, see Appendix A.

Architecture. Figure 1 shows our measurement setup. We host a web server and a corresponding authoritative name server (denoted as “a.com”, located in the U.S.) to receive DNS and HTTP requests. The authoritative name server runs BIND9 on Linux [27]. We also control a measurement client to communicate with the

BrightData Super Proxy, which instructs exit nodes to resolve our domain name either via DoH or Do53. BrightData does not allow our measurement client to directly control the exit node, rather, all requests must be routed via the BrightData Super Proxy.

We use BrightData not only for its global coverage of exit nodes, but also for a number of features that facilitate our measurements: (1) We can specify the *country* of an exit node for a particular request, allowing us to target clients globally. (2) We can make multiple requests via the same exit node. This allows us to measure both DoH and Do53 performance from a single exit node.

Measurement Workflow. Our measurement client takes a *country code* and *target public DoH resolvers* as input. After a survey of relevant literature [15, 29], we selected four public DoH providers to examine for this study: Cloudflare [13], Google [17], NextDNS [37], and Quad9 [40]. These servers include some of the largest DoH providers and we view them as representative of current public DoH offerings. Our client first connects to the BrightData Super Proxy and requests to connect to an exit node in the specified country. The Super Proxy randomly selects an exit node in the given country and then acts as the middle-man to forward our traffic to the exit node. For each exit node, we run two distinct measurements:

- **DoH Measurement** The exit node performs a DoH resolution for a unique subdomain of our web server (e.g., <UUID>.a.com) for each public DoH resolver. We use a unique subdomain (e.g., a UUID) for each request to control for any domain caching issues, thereby forcing the client to contact our authoritative name server for each measurement. In the process, the public DoH resolver queries our authoritative name server (as shown in Figure 1, steps 3–6). Note that the Super Proxy itself does not implement any code to perform DoH resolutions. We send the DoH resolution request from our measurement client and the Super Proxy only acts as the middle-man to forward the request to the exit node.
- **Do53 Measurement** To conduct a Do53 measurement, the exit node sends an HTTP GET request to our web server at a unique subdomain (again, <UUID>.a.com). This triggers a Do53 resolution at the exit node. We note that the exit nodes may be configured to use a variety of DNS resolvers (i.e., from ISP-provided resolvers to custom resolvers). This methodology allows us to measure the Do53 performance under each individual exit node’s *default* configuration. This assumption is verified in Section 4.3

We made a conscious decision to control the impact of DNS caching for all our measurements. By using a unique subdomain (e.g., a UUID) for each request, we force the client to contact our authoritative name server each time. The purpose is to rule out the impact of caching while allowing us to attribute differences in resolution time to the transport protocol instead of the resolved domain name. This approach is similar to that used in prior works [8, 29]. Although this method does not capture clients’ cache hit performance, it represent a “worst-case” evaluation for both Do53 and DoH. We will further discuss this limitation later in Section 7.

3.2 Calculating DoH Query Time

As noted earlier, we do not have a full control over the exit nodes to *directly* run measurements for the DoH and Do53 resolution times.




Figure 2: DNS-over-HTTPS (DoH) Request Timeline—Our measurement client sends a DoH request to the BrightData SuperProxy, which forwards the request to an exit node. The exit node then sends a DoH request to an DoH provider (e.g., Cloudflare), which resolves the domain name by contacting our authoritative name server.

In this section, we detail our strategies to measure, derive, and verify the resolution time based on the timing information collected from our measurement client, our web server, our authoritative name server, and the information obtained from the Super Proxy.

Figure 2 shows a detailed breakdown of the measurement process for DoH (22 total steps). We use t_i to denote the time taken in the i^{th} step. For each measurement, we instruct the exit node to resolve a unique subdomain name “<UUID>.a.com” under our control by sending a HTTPS request to a public DoH resolver (e.g., cloudflare.com).

Step (1–8): Establish a TCP Connection to DoH Server. Steps (1–2) initiate the establishment of a TCP tunnel (using HTTP CONNECT) from our measurement client to the exit node via the Super Proxy.

In step (3–4), the exit node first resolves the DoH server’s domain name (e.g., cloudflare.com) with its local DNS configuration. After that, the exit node does a 3-way TCP handshake with the DoH server (steps 5–6).

In step (7), the exit node replies back to the Super Proxy. The HTTP response headers contain useful timing information, for example, the time it took to resolve the domain name of the DoH server ($t_3 + t_4$) and the time of the TCP handshake ($t_5 + t_6$). In step (8), the Super Proxy sends a “200 OK” back to our client, establishing the TCP tunnel. In addition, it sends our measurement client the timing information encoded in the response headers from the Super Proxy.

Step (9–14): Establish a TLS Session. Using the TCP tunnel, our client sends a ClientHello to establish a TLS session with DoH resolver in step (9–11). The DoH resolver then sends back a ServerHello and Finished to our client in step (12–14). Note that we establish a TLS session in only one round trip due to the use of TLS 1.3 as specified in RFC 8446 [42] (TLS 1.2 uses two round trips). Since TLS 1.3 is now supported and preferred in the DoH resolvers we study, we only include one round trip time.

Step (15–22): Resolve the Domain Name. In step (15–17), our client sends a Finished and an HTTP GET request to the DoH resolver to resolve the target domain name “<UUID>.a.com”. The DoH resolver then resolves the domain name by contacting our authoritative name server in step (18–19). In step (20–22), the DoH server encrypts the resolved IP and sends it back to our client, completing the DoH resolution.

3.2.1 Calculating DoH Resolution Time. Our goal is to measure the round trip time for DoH resolving at the exit node. To mimic reality, we need to exclude the time spent to communicate with the

Super Proxy, and thus the total time is:

$$t_{DoH} = (t_3 + t_4 + t_5 + t_6) + (t_{11} + t_{12}) + (t_{17} + t_{18} + t_{19} + t_{20}) \quad (1)$$

Known Timing Information. To calculate t_{DoH} , we rely on three sets of available timing information. First, on the measurement client, we can obtain four timestamps (marked out as A, B, C, and D in Figure 2).

Second, we can calculate $(t_3 + t_4 + t_5 + t_6)$ based on the header information from BrightData’s Super Proxy. Specifically, BrightData’s Super Proxy collects important timing information from the exit node. In the HTTP header (received at step 8), the X-luminati-tun-timeline field has two key values: the “DNS” value is $t_3 + t_4$, and the “Connect” value is $t_5 + t_6$.

Third, we can obtain the processing time spent on BrightData boxes (Super Proxy and exit node), denoted as $t_{\text{BrightData}}$. This is done based on the HTTP header from the Super Proxy. Header field X-luminati-timeline includes the detailed time spent on BrightData boxes to authenticate the client, initialize the Super Proxy, select and initialize the exit node, and check the validity of the requested domain name. We obtain $t_{\text{BrightData}}$ by simply adding the provided times.

Assumptions. Our remaining calculation is based on two assumptions. We will validate the assumptions and our overall methodology in a ground-truth experiment detailed in Section 4.

- (1) We assume the round trip time between our client and the exit node is relatively stable. This means $RTT = (t_1 + t_2 + t_7 + t_8) = (t_9 + t_{10} + t_{13} + t_{14}) = (t_{15} + t_{16} + t_{21} + t_{22})$.
- (2) The processing time spent by BrightData boxes ($t_{\text{BrightData}}$) is only incurred once when we establish the TCP tunnel (step 1–8). Once the tunnel is established, BrightData boxes take negligible time to forward later requests. (step 9–22).

Calculating t_{DoH} . With these assumptions, we can now calculate t_{DoH} based on Equation 1.

First, in Equation 1, $(t_3 + t_4 + t_5 + t_6)$ is already provided by the Super Proxy, and we only need to calculate the remaining parts. Based on the two timestamps T_C and T_D , we have:

$$T_D - T_C = \sum_{i=9}^{22} t_i \quad (2)$$

The above equation involves Assumption-2 as we assume BrightData boxes take minimal time to forward the request after the initial TCP connection is established. Then based on Assumption-1, the round trip time between our client and exit node (RTT) stays the

same for $(t_9 + t_{10} + t_{13} + t_{14})$ and $(t_{15} + t_{16} + t_{21} + t_{22})$, and thus the above Equation 2 can be rewritten as:

$$t_{11} + t_{12} + t_{17} + t_{18} + t_{19} + t_{20} = T_D - T_C - 2 \times RTT \quad (3)$$

Then, by adding $(t_3 + t_4 + t_5 + t_6)$ to both sides, we have:

$$t_{DoH} = T_D - T_C + (t_3 + t_4 + t_5 + t_6) - 2 \times RTT \quad (4)$$

At this point, we only need to calculate RTT to obtain the desired t_{DoH} . To calculate RTT , we use the two timestamps T_A and T_B , and compute:

$$T_B - T_A = \sum_{i=1}^8 t_i + t_{\text{BrightData}} \quad (5)$$

where $t_{\text{BrightData}}$ is the time spent on BrightData boxes to establish the TCP tunnel (already known). As stated in Assumption-1, the round trip time stays the same as $RTT = t_1 + t_2 + t_7 + t_8$. We can rewrite the above equation as:

$$RTT = T_B - T_A - (t_3 + t_4 + t_5 + t_6) - t_{\text{BrightData}} \quad (6)$$

By taking Equation 6 to Equation 4, we have:

$$t_{DoH} = (T_D - T_C) - 2 \times (T_B - T_A) + 3 \times (t_3 + t_4 + t_5 + t_6) + 2 \times t_{\text{BrightData}} \quad (7)$$

We obtain t_{DoH} based on Equation 7 where all the values are known from measurements/header information.

3.3 Calculating Do53 Query Time

For Do53 measurements, we simply extract the timing information from the header of BrightData's response. Recall that our Do53 measurement is to instruct the exit node to visit our website under `<UUID>.a.com` via the Super Proxy. During the process, the exit node uses traditional DNS resolving (e.g., DNS-over-UDP) with their default configurations. The query time of Do53 is recorded in the header of Super Proxy response (`X-luminati-tun-timeline` header; "DNS" value). We validate the reliability of the Super Proxy's header information in Section 4.2.

3.4 DoH Connection Reuse

Existing studies show that DoH performance can be improved if a user reuses the same TLS connection for multiple DNS resolutions [8, 22]. As such, we also want to measure the performance of DoH connection reuse. We denote t_{DoHR} as the DoH query time if the exit node reuses an already established TLS session to send more DNS queries. In this case, t_{DoHR} represents the performance of subsequent queries (after the first query). t_{DoHR} is expected to be shorter than t_{DoH} as we no longer need to perform a TCP handshake or TLS session establishment.

Directly measuring DoH connection reuse is not feasible at the exit node. This is because the BrightData Super Proxy closes connections after a request is sent. To estimate t_{DoHR} , we calculate an upper bound value by subtracting the time for DNS resolution ($t_3 + t_4$), TCP handshake ($t_5 + t_6$) and TLS session establishment ($t_{11} + t_{12}$) from t_{DoH} . This means, $t_{DoHR} = t_{DoH} - (t_3 + t_4 + t_5 + t_6) - (t_{11} + t_{12})$. Based on Equation 7, we have:

$$t_{DoHR} = (T_D - T_C) - 2 \times (T_B - T_A) + 2 \times (t_3 + t_4 + t_5 + t_6) + 2 \times t_{\text{BrightData}} - (t_{11} + t_{12}) \quad (8)$$

In this equation, all the values are known so far except for $(t_{11} + t_{12})$. To obtain $(t_{11} + t_{12})$, we assume the round trip time between the exit node and the DoH resolver is near identical, which means $(t_{11} + t_{12}) = (t_5 + t_6)$ (see Figure 2). As stated before, $(t_5 + t_6)$ is known based on the Super Proxy's header. With this assumption, all the values in Equation 8 are known and we can estimate t_{DoHR} . We validate this calculation method in Section 4.1.

3.5 Limitations and Remedies

Our measurement methodology works for the vast majority of the countries in the BrightData network, with a few exceptions. These exceptions only apply to Do53 measurements and do not affect DoH data. Specifically, we find that our method (Section 3.3) cannot return accurate Do53 measurements for 11 countries (out of 200+ countries, 5%) where the BrightData Super Proxy servers are located [5]. These 11 countries include the USA, Canada, UK, India, Japan, South Korea, Singapore, Germany, Netherlands, France, and Australia. In these 11 countries, BrightData will perform DNS resolution at the Super Proxy rather than at the exit node regardless of our request configuration. As a result, the header information we obtain does not reflect the Do53 query time at the exit node.

To obtain the missing Do53 data from the 11 countries listed above, we leverage the RIPE Atlas network [43]. RIPE Atlas is a global volunteer network to support simple connectivity and reachability measurements. RIPE Atlas supports conventional DNS probing, which is sufficient to collect Do53 data from those 11 countries.²

To ensure the Do53 measurement data obtained from RIPE Atlas is consistent with the rest the data from BrightData, we perform validation experiments in "overlapping" countries that are covered by both BrightData and RIPE Atlas (see Section 4.4). The experiments confirm that our remedy strategy is valid. Given that RIPE Atlas does not support DoH measurements, for these 11 countries, we combine the Do53 data from RIPE Atlas with the DoH data from BrightData for our analysis. The Do53 data from the 11 countries can support most of our analyses, with the exception of any per-client DoH-Do53 comparisons (Section 6).

Another limitation of BrightData is that it might not always be accurate in mapping exit nodes to the country the user resides in. Based on our observation, BrightData uses the IP Address of an exit node to determine the country it is from but there might be a chance that BrightData makes mistakes. To account for these inaccuracies, we add an additional check on our end. As mentioned in Section 3.1, for Do53 measurements, the exit node sends an HTTP GET request to our web server. Thus we know the /24 subnet prefix of the exit node. We use this prefix to determine the location of the exit node using the Maxmind Geolocation Service [32]. We discard any data points for which there is a mismatch between the country specified in the BrightData API and the country determined using the Maxmind service. In the end, we discarded 0.88% of the data

²RIPE Atlas does not support HTTPS connections to arbitrary hosts and thus we do not use RIPE Atlas for DoH measurements in the first place.

Country Query Time (ms)	Ireland		Brazil		Sweden		Italy		India		USA	
	DoH	DoHR	DoH	DoHR	DoH	DoHR	DoH	DoHR	DoH	DoHR	DoH	DoHR
Our Method	116	94	193	182	129	122	246	236	254	251	53	25
Ground-Truth	109	85	190	176	131	126	245	238	260	257	52	23
Difference	7	9	3	6	2	4	1	2	6	6	1	2

Table 1: Ground-truth Experiments for DoH and DoHR—We set up our own exit nodes in 6 different countries to valid our DoH measurement methodology. We show the median DoH and DoHR query time (in millisecond) obtained by our method match well with the ground-truth.

Country	Ireland	Brazil	Sweden	Italy
Our Method	102	139	131	204
Ground-Truth	102	138	129	203
Difference	0	1	2	1

Table 2: Ground-truth Experiments for Do53— We set up our own exit nodes in 4 different countries to valid our DoH53 measurement methodology. We show the median query time (in millisecond). Do53 measurement is not applicable via BrightData in the USA and India (see Section 3.5).

points collected. All the analysis and results in the rest of the paper have already excluded such data points.

4 GROUND-TRUTH VALIDATION

Our measurement methodology is based on several assumptions outlined in Section 3. In this section, we run small-scale experiments to test the validity of our methodology before running full measurements.

4.1 Validating DoH and Connection Reuse

To validate the correctness of our DoH measurement methodology (Section 3.2 and Section 3.4) we set up our own machines in different locations and volunteer them to join the BrightData network as exit nodes. We then force the BrightData Super Proxy to select our own machines as exit nodes to perform tests. Once our measurement client is successfully connected to our own exit node (via the Super Proxy), we are able to perform a “ground-truth” DoH measurement at the exit node and compare it with the values calculated by our proposed method.

Setup. We set up six EC2 machines (with full control) in Ireland, Brazil, Sweden, Italy, India, and the USA. For each machine, we install the HolaVPN software to make the machine part of the BrightData network of exit nodes. We then repeatedly query the Super Proxy with the corresponding country code, city name, and ASN of our machine until our machine is eventually selected as the exit node. Because we have a full control over the measurement client and the exit node, we can then obtain a complete view of how the BrightData network works.

DoH Validation Experiment. To validate our DoH measurement method (Equation 7), we first directly control the exit node to perform a DoH resolution with a DoH resolver (i.e., Cloudflare) and record the query time, which we consider “ground-truth”. Then, we run our proposed DoH measurement method to obtain the query time—in this case, the DoH query is performed via the SuperProxy,

and the DoH query time is calculated using Equation 7. For each machine, we repeat measurements 10 times and take the median query time. Table 1 shows exact and estimated DoH measurement times for each ground truth node. Our method returns consistent values compared to ground-truth measurements, with differences within 8ms.

DoH Connection Reuse. We also validate our method to calculate query timing for DoH connection reuse t_{DoHR} (Section 3.4, Equation 8). The ground-truth t_{DoHR} is obtained by directly controlling the exit node to perform a DoH query multiple times and re-using the same TLS connection. As shown in Table 1, the DoHR query time obtained by our method is highly consistent with the ground-truth values across all six countries, again with differences under 10ms per query.

4.2 Validating Do53 Measurements

We conduct similar validation experiments for our Do53 measurements. Recall that USA and India are among the 11 countries that have BrightData Super Proxy servers and thus the Do53 measurement is not applicable (see Section 3.5). As such, we only run the Do53 validation experiments on the other 4 machines. We compare the time taken to conduct a Do53 measurement at each ground truth node with the Do53 query time collected from the Super Proxy header (see Section 3.3). For each machine, we repeat this experiment 10 times to report the median value. We find that our Do53 measurement method is consistent, with differences within 2ms (Table 2).

4.3 Default DNS Protocols of Exit Nodes

Our methodology assumes the default DNS resolving protocol of the exit nodes is Do53. Here we briefly justify and validate this assumption. When we proxy a request through the exit nodes, there could be different possible ways for the exit nodes to resolve the domain names. They could be using browser specified configurations (since HolaVPN is a browser extension), the default operating system settings, or even specific DNS servers hard coded in the HolaVPN. To figure out the DNS resolution mechanism used by exit nodes, we perform experiments using these exit nodes under our control. We instruct each exit node to visit our website under “<UUID>.a.com” via the Super Proxy multiple times. For each measurement, we use different resolvers configured for the operating system and the browser. We then capture the packets on our machines using Wireshark. In these experiments, we observe that all of the exit nodes consistently use the resolver configured for the operating system as the default.

At the time of our measurement, no major operating system (e.g., Windows, MacOS, Linux, Android, iOS) is configured to use DoH by default [1, 7, 12, 31, 45], and thus it is unlikely that clients of BrightData have configured their operating systems to use DoH. As such, we assume Do53 is still their default configuration. It is certainly possible that some ‘tech-savvy’ users may have changed their own settings to use DoH by default at the operating system level or at the network level. However, analysis of our measurement data (Section 5) provides further support for our assumption, given the significant discrepancies between the DoH measurement results and those of the default resolvers. This at least suggests that such tech-savvy users are uncommon.

4.4 BrightData vs. RIPE Atlas

Finally, we validate our remedy strategy that will be used to collect the missing Do53 data from the 11 countries where BrightData Super Proxy servers are located (Section 3.5). In this experiment, we examine whether BrightData and RIPE Atlas return consistent Do53 measurements in a given country, which is important to decide whether we can combine the Do53 data collected from the two networks in our analysis.

To do this, we randomly select 10 countries where both BrightData and RIPE Atlas can obtain valid Do53 measurements.³ For each country, we run the Do53 measurements at both networks at least 250 times and take the median Do53 query time. We find that results from the two networks are highly consistent. Across the 10 countries, the average difference in Do53 query time between the two networks is only 7.6ms with a standard deviation of 5.2ms.

Summary. Our ground truth experiments suggest that our methodology is sound and can be used to closely approximate actual-valued DoH and Do53 resolution times around the world.

5 MEASUREMENTS

In this section, we characterize our dataset and present initial findings across four major DoH providers. For each provider, we examine the time it takes for clients in various countries to perform DNS resolutions using both DoH and their default resolvers.

Terminology. In the remainder of the paper, we refer to DoH measurements in two ways. The first, as described in Equation 8, is *DoHR*, which refers to the time it takes to complete a DoH request after the TLS connection has already been established (to emulate connection-reuse). The second is *DoHN*, where *N* is an integer denoting the number of requests made over a single connection. This notation expresses the average resolution time over *N* requests, beginning with the TLS handshake on the first request. For example, *DoH1* describes the time it takes for a single DoH request, including the initial handshake. *DoH10* measures the average per-request time over a DoH connection that handles 10 resolutions.

5.1 Dataset

Our dataset was collected during April and May of 2021. It consists of 22,052 unique clients from 224 unique countries and territories. We limit our per-country analyses to countries where we are able

³These 10 selected countries include Belgium, South Africa, Sweden, Italy, Iran, Greece, Switzerland, Spain, Norway, Denmark, New Zealand, Austria, and Bulgaria.




Figure 3: Clients Per Country in Dataset—The distribution of clients per country in our dataset for those countries included in our per-country analysis. In the median case, we analyze 103 unique clients per country, but we have at least 200 clients for 17% of countries .

Resolver	Clients	Countries
Cloudflare	21,858	222
Google	21,905	223
NextDNS	21,947	223
Quad9	21,897	223
Do53 (Default)	22,052	224

Table 3: Dataset Composition—The distribution of our data for each resolver. For all four DoH resolvers, we have data points for at least 21,858 unique clients spanning at least 222 unique countries.

to obtain at least 10 unique clients that perform a resolution using each of four selected DoH providers (Cloudflare, Google, NextDNS, Quad9). This causes us to exclude 25 countries and territories including China, North Korea, Saudi Arabia, and Oman. Our clients are located in 2,190 different autonomous systems (ASes), and we observe queries to our authoritative DNS server from 1,896 unique recursive resolvers. Table 3 shows a breakdown of our clients across different resolvers and countries, and Figure 3 shows the number of clients per country across our dataset. We apply the Maxmind Geolocation Service [32] to get approximate latitude and longitude using the /24 for each client as well as the true IP address for each recursive resolver which we observed querying our authoritative name server. For each client, we send 5 total requests in one measuring run: One to measure each of the four DoH providers we study, and one to the default DNS resolver for that client. We conduct 2 runs per client. While public DNS providers often offer special services for malware and adult content blocking, we simply use the default resolution service for each public DoH provider.

5.2 Differences Between DoH Providers

Figure 4 shows CDFs of the client resolution times for *Do53* as well as DoH with and without the initial TLS handshake. Then, in Figure 5, we examine the performance of each provider geographically. Finally, by geolocating clients and resolvers, we investigate whether real-world clients are using the closest available points-of-presence (PoPs) for each provider. We approximate the “potential improvement” for a client as the difference between the distance from the client to the DoH PoP it actually used, and the distance from the client to the closest PoP (of the same DoH provider) in our dataset. Figure 6 shows this distribution for each provider. We




Figure 4: Resolution Times by Resolver—We show the distribution of resolution times for each public DoH resolver, both for the initial (*DoH1*) and repeated (*DoHR*) requests, along with the default (*Do53*) distribution for reference. Cloudflare stands out, as its *DoHR* resolution times very closely track *Do53* times for most clients.

note that geographic distance is an imperfect proxy for network distance and latency, however, the results help to explain some of the differences we observe between providers. Below, we discuss the results for each resolver.

Cloudflare. Cloudflare is the top-performing DoH provider in our study with a median *DoH1* resolution time of 338ms. Figure 4(a) shows that after a TLS handshake is completed, DoH resolution via Cloudflare performs similarly to Do53 through a client’s default resolver. In the median case, each subsequent DoH request (*DoHR*) through Cloudflare takes 257ms, compared to a median of 250ms for Do53 queries.

From a geographic perspective, Cloudflare also performs well. As shown in Figure 5(a), we observed 146 unique PoPs across the globe for Cloudflare—the most out of the four providers we studied. In addition to providing the best DoH speeds in many western countries, Cloudflare’s large number of points of presence help it provide respectable speeds in regions where other providers struggle. For example, Cloudflare is the only provider with a PoP in Senegal, and the median resolution speeds for Cloudflare (274ms) are significantly better than the next provider, Google (381ms).

Google. The median *DoH1* resolution time for clients using Google is 429ms—the second best of the four providers. However, once that initial connection is established and we measure the timing of subsequent requests (*DoHR*), Google (315ms) falls behind Quad9 (298ms) to third place overall in median resolution time.

Google’s map in Figure 5(b) stands out due to the lack of points of presence compared to the other providers. We observed only 26 unique PoPs for Google, not finding a single one in Africa. The Google Cloud website [11] states that Google employs at least 61 PoPs, including 3 in Africa, but it is unclear whether all of these PoPs provide DoH service (some PoPs may be set up to provide other services such as content delivery). This raises the question—How does Google provide resolution times on-par with other providers with fewer PoPs? Figure 6 offers a possible explanation. Google appears to minimize the number of clients who are using unnecessarily distant PoPs relative to other providers. Only 10% of Google clients could be switched to a PoP at least 1000 miles closer, as compared to 26% of Cloudflare clients.

NextDNS. NextDNS does not have its own autonomous system (AS) on which its resolvers operate, which is different from the other providers we examined. Instead, NextDNS’s 107 points of

presence we observed are achieved through recursive resolvers distributed across at least 47 different ASes. Interestingly, these ASes include Google and Cloudflare, indicating that NextDNS may be routing queries through those two companies in a limited number of cases. Perhaps because of this, NextDNS has the slowest DoH performance in our dataset, both relative to Do53 (1.47x) performance and overall (median *DoH1* of 467ms). Despite this, however, NextDNS DoH performs extremely well in the United States, bested only by Cloudflare (206ms vs 163ms).

Quad9. We find Quad9 to be in the middle of the pack performance-wise, with a median *DoH1* time of 447ms and adding 28% overhead relative to *Do53* over 10 requests (*DoH10*). Looking at the map for Quad9, shown in Figure 5(d), we see that Quad9 has far more points of presence in Sub-Saharan Africa than other resolvers, but this does not seem to provide obvious benefit, as Quad9 performs similarly to Google and Cloudflare in these regions. We find that Quad9 has significant room for improvement in the way it assigns clients to geographic PoPs relative to other providers (Figure 6). For the median Quad9 client in our dataset, there is a PoP 769 miles closer than the one that was used. Although geographic distance does not necessarily reflect network distance or latency, median potential improvement for Quad9 is significantly worse than Cloudflare (46 miles), Google (44 miles), and NextDNS (6 miles), highlighting Quad9 resolvers as distinct outliers.

5.3 Geographic Differences

In addition to measuring DoH across different resolvers, we also study differences in aggregate DoH performance across countries, finding significant variation. In the median case, countries had a *DoH1* time of 564.7ms and a *Do53* time of 332.9ms. However, in both cases, clients from several countries took significantly longer. For example, clients from Chad, a country in Central Africa, took 2011ms to resolve our initial DoH queries and 1280ms to resolve our *Do53* queries. In contrast, clients from some countries have significantly faster DoH and *Do53* resolution times, for example, Bermuda, which has a median *DoH1* time of 204.1ms and a median *Do53* time of 90.5ms.

We compare our DoH and *Do53* measurements within countries by calculating the delta between the medians of country-wide resolution times. Figure 7 shows the delta aggregate per country, split by the resolver used to complete the initial DoH request. For




Figure 5: DNS Resolution Times and Points of Presence (PoP)—We show the median *DoH10* resolution time for each country in our dataset. Points of presence (PoP) we observed for each provider are shown as black stars. The greenest country (NextDNS-Canada) has a median resolution time of 63ms, while the reddest nations have median resolution times of over 1 second. The same color scale is consistently used across the four maps. A small number of countries and territories, most notably China, remain gray as we were unable to obtain DoH resolution data across all four public providers for them.




Figure 6: Potential Improvement in Distance to DoH PoP—We define “potential improvement” as the difference between the distance from the client to the DoH PoP it actually used, and the distance from the client to the closest PoP in our dataset. Although Google has fewer PoPs than other providers, it assigns a higher percentage of clients to the closest PoP, compared to Quad9, who appears to have significant room for improvement.

most countries, a switch to DoH increases the time taken to perform a single DNS query, which is expected. However, we note that for 8.8% of countries, switching to DoH actually reduces the time taken to perform a single DNS query. For example, clients in

Brazil experienced a 33% speedup in DNS performance with *DoH1* compared to *Do53*. Although we cannot say conclusively why this happens based on our data, we provide more information on the types of countries where this occurs in Section 6. Similar to our client-centric results, DoH resolutions from Cloudflare cause the smallest performance hit by this metric, with the median country experiencing a relatively modest (19%) performance decrease compared to resolvers from Quad9, Google, and NextDNS, who cause a 28%, 39%, and 47%, and performance decrease per country respectively.

6 EXPLAINING DIFFERENCES IN DOH PERFORMANCE

In this section, we identify country- and client- level factors that can explain DoH and Do53 performance differences between clients.

6.1 Identifying Explanatory Variables

Our main questions are to understand if countries with developing economies and developing Internet infrastructure are disproportionately impacted by a universal switch to DoH. As a proxy for economic development in a country, we collect Gross Domestic




Figure 7: DNS Performance by DoH Resolver—We show the change in resolution times when switching from *Do53* to *DoH10*. The change differs significantly across DoH providers, with Cloudflare causing a slowdown of 49.65ms in the median case, while NextDNS causes a significantly greater slowdown of 159.62ms.

Variable	OR	OR_10	OR_100	OR_1000
Bandwidth (Control = Fast)				
Slow	1.81x	1.69x	1.66x	1.65x
Income Group (Control = High)				
Upper-middle	1.50x	1.06x	1.00x	0.99x
Lower-middle	1.76x	1.27x	1.20x	1.19x
Low	1.98x	1.37x	1.27x	1.25x
Num ASes (Control = Higher than Med)				
Lower than Med	1.99x	1.76x	1.70x	1.69x
Resolver (Control = Cloudflare)				
Google	1.76x	1.77x	1.71x	1.70x
NextDNS	2.25x	1.99x	1.91x	1.90x
Quad9	1.78x	1.34x	1.27x	1.25x

Table 4: Modeling DoH vs. Do53 Slowdowns—We show the results of our logistic regression with the categorical variable inputs. All results are statistically significant with $p < 0.001$.

Product (GDP) per capita data from the World Bank [3]. To proxy for Internet infrastructure development, we leverage data collected by Ookla’s Speedtest service which publishes a global index of fixed broadband speeds per country [39] and collected the number of ASes per country collected by IPInfo [26].

6.2 Modeling DNS Performance

In order to account for the impact that multiple latent variables may have on our performance measurements, we rely on regression analysis with each of our explanatory variables as inputs. Regression analysis enables us to measure the impact of a single variable while holding all other variables constant. We developed two models to investigate 1) why countries might experience worse-than-median DoH slowdowns, and 2) the strength of each explanatory variable in relation to one another.

6.2.1 Logistic Modeling. We model the outcome of the transition from Do53 to DoH as a *multiplier* between Do53 and DoH1, DoH10, DoH100, and DoH1000 respectively. We convert the multiplier to a binary outcome based on if the multiplier is better or worse than the global median, which is 1.84x, 1.24x, 1.18x, and 1.17x for 1, 10, 100, and 1000 requests respectively. We treat clients that achieve a

multiplier lower than this value as a successful event (or a 1), and clients that achieve a multiplier greater than this value as a failure (or a 0). We leverage four categorical input variables:

- (1) **Bandwidth.** One of “Fast” or “Not fast”. Determined by the United States Federal Trade Commission’s definition of “fast Internet speed” ($> 25\text{Mbps}$) [14]
- (2) **Income Group.** One of “High income”, “Upper middle income”, “Lower middle income”, or “Low income”. Determined via GDP data by the World Bank [3].
- (3) **Number of ASNs** One of “High” or “Low”. Determined by if a country had higher than the median number of ASNs per country globally (25 ASes).
- (4) **Resolver.** One of “Cloudflare”, “Google”, “NextDNS”, or “Quad9”.

Table 4 shows the results of our logistic regression. We report effect sizes as the odds that a client with a particular property—after holding all other features constant—will experience a speedup or slowdown when transitioning to *DoHN* from *Do53*. We detail our results for each feature:

Bandwidth. We find that the odds of experiencing a slowdown when transitioning to DoH from Do53 is 1.81x for clients with slow Internet connections compared to those with fast Internet connections for a single request. This trend does not significantly change even when the TLS tunnel is reused for multiple DNS requests—even if a single connection was used for 1000 queries, the odds of a client with low bandwidth experiencing a slowdown are 1.65x compared to clients with fast Internet speeds. Clients with low Internet speeds experience a median slowdown time of 350ms for a single request, compared to a median slowdown time of 112ms for clients with fast Internet speeds, approximately a 3.1x slowdown.

Income Group. The odds that clients from low income countries experiencing a slowdown is 1.98x compared to clients from high income countries. We also observe a direct, linear relationship between income levels of countries and the odds that their clients will experience a slowdown—the odds clients from lower-middle income countries experience a slowdown is 1.76x compared to clients from high-income countries; the odds that clients from upper-middle countries experience a slowdown is 1.5x compared to clients from high-income countries. The trend is significantly dampened, however, when considering multiple requests—if a single connection was used to perform just 10 DNS queries, the odds that low-income countries experience a slowdown is reduced to just 1.37x, indeed highlighting the benefits that using a single TLS session can afford to countries with varied income groups.

Despite these relative improvements, clients from low-income countries still experience a significant raw slowdown—the median slowdown is 461ms compared to a slowdown of 84ms for high-income countries for *DoH1*. Although raw performance improves for clients from middle income countries with additional requests (just a 52ms slowdown at *DoH100*), we do not observe a similarly scaled difference for clients from low-income countries, who experience a median 200ms slowdown at the 100th request.

Number of ASes. The number of ASes in a country plays a similar role as bandwidth as an explanatory factor, as both are proxies for Internet infrastructure investment. As such, we see the

Output	Metric	Coef. (ms)	Scaled Coef. (ms)
Delta	GDP	-6.67e-4*	-13.8*
	Bandwidth	-2.26	-134.5
	Num ASes	-5.9e-2	-80.8
	Nameserver Dist.	1.13e-2	30.0
	Resolver Dist.	5.6e-2	93.4
Delta 10	GDP	-3.5e-4*	-7.3*
	Bandwidth	-1.23	-73.3
	Num ASes	-4.7e-2	-63.6
	Nameserver Dist.	7.3e-3	19.6
	Resolver Dist.	2.6e-2	42.4
Delta 100	GDP	-3.2e-4*	-6.6*
	Bandwidth	-1.13	-67.2
	Num ASes	-4.6e-2	-61.9
	Nameserver Dist.	7.0e-3	18.5
	Resolver Dist.	2.3e-2	37.3

Table 5: Linear Modeling of DNS Performance—We show the results of our linear modeling, with both unscaled and scaled coefficient values for maximal interpretability. Internet infrastructure investment is the most significant factor to consider when evaluating DoH performance slowdowns worldwide. All results are statistically significant with $p < 0.001$ with the exception of GDP, which was not significant for any regressions.

respective odds ratios follow a similar pattern—the odds that countries with a lower number of ASes than the median experience a slowdown is 1.99x compared to countries with a higher number of ASes than the median. There is a slight change when additional DNS requests are added, however, the number of ASes in a country is still a strong predictor of whether clients in a given country will experience a slowdown. Although some studies have suggested DoH can improve performance compared to Do53 [22], we observe this effect is rarely attributed to clients with low Internet infrastructure investment—of the clients that experience a DoH *speedup* compared to Do53, 84% have fast nationwide Internet speeds and 93% have high numbers of ASes.

Resolver. The choice of DoH resolver has a significant impact on client performance. For almost all requests, we observe that clients that used Google, NextDNS, and Quad9 for DoH resolutions experienced a significant slowdown compared to clients that used Cloudflare for DoH resolution—an odds of experiencing a slowdown of 1.76x, 2.25x, and 1.78x respectively. These odds decrease slightly as more requests are added, the largest increase of which goes to clients that use Quad9, for which the odds increase only 1.27x compared to Cloudflare for the 100th request. In spite of these relative differences, we observe that in aggregate, the raw slowdowns experienced by each resolver are relatively similar, ranging from 28ms to 85ms for *DoH100*, indicating that the performance impact from choice of resolver is significantly reduced when considering multiple requests for the same TLS session.

Summary. Our results indicate that a universal transition to DoH from Do53 would disproportionately impact countries with lower income and less Internet infrastructure investment.

6.2.2 Linear Modeling. Although our logistic model gives us insight into which types of clients will experience a slowdown when

transitioning to DoH, it does not tell us the impact that each explanatory variable has on the *continuous* outcome of the delta between Do53 and DoH times. To measure this, we model the *raw time delta* between Do53 and DoH per client. We model this outcome as a linear regression with the GDP per capita, broadband speed, number of ASes available in each country, the geodesic distance from the client to our authoritative name server, and the geodesic distance from the client to the DoH resolver used as input variables. We add distance metrics as they may serve as latent confounds to other results—by controlling for distance, we remove concerns about bias introduced because clients were closer or further away from the DNS infrastructure used to resolve the query.

Table 5 shows the results of our model. We report the model weights as coefficients of the linear regression, which shows the relative impact of each individual variable on the outcome. We also show normalized coefficients, which is the outcome of the linear regression after scaling each explanatory variable to a scale from 0 to 1. All results presented were statistically significant with $p < 0.001$ with the exception of GDP, for which no results were statistically significant.

For each output, we observe that a client country’s investment in Internet infrastructure is the strongest predictor of a DoH slowdown. For *DoH1*, a difference of 1Mbps in nationwide bandwidth has an estimated impact of -2.26ms. Normalized, a change in one unit of nationwide bandwidth or the number of ASes has an estimated impact of -134.5ms and -80.8ms in delta performance respectively. We did observe a small trend that the distance a client is from our authoritative nameserver increases the delta time, but this is far outstripped by Internet investment factors. In contrast, the second largest factor in predicting delta time was the distance to DoH recursive resolver—one normalized unit change in resolver distance amounted in an estimated impact of 93.4ms in query time. As noted in Section 5, DoH providers have different PoP placement strategies, with Cloudflare opting for a more globally distributed presence while Google tends to have smaller, more centralized PoPs that handle more geographic area. Even when considering clients from a single DoH resolver (e.g., Google), the distance between the recursive resolver and the client has a statistically significant impact on delta performance, even matching Internet investment features for Google and Cloudflare. These results highlight that resolver deployment strategy and efficient routing will play an important role in equitable DoH performance. We show full resolver-filtered regression tables in Appendix C.

Increasing the number of DNS requests per TLS connection decreases the scale of each coefficient, and notably decreases the relative power that bandwidth has when compared to the number of ASes per country (1.7x to 1.1x), noting that bandwidth may play a smaller role in practice than nationwide Internet investment broadly. Both Internet investment features outweigh distance and nationwide income metrics when multiple requests are considered. Our results highlight that as we move towards deploying DoH universally, we should consider the impact that the protocol will have on Internet clients worldwide and potentially change our deployment strategies to not disproportionately affect clients with lesser means.

7 DISCUSSION AND FUTURE WORK

Our results point towards solutions for a more equitable DoH-by-default deployment, with implications for both software vendors (e.g., browsers, operating systems) and DoH resolution services.

Software Vendors. Countries with already low Internet infrastructure investment and economic development will be disproportionately impacted by a unilateral switch to DoH from Do53. As such, we suggest that software vendors refrain making DoH the *default* choice for DNS resolution for clients, at the least until measurement data for each country suggests that the impact of turning on DoH would be negligible for common Internet applications. In some cases, the performance cost may be acceptable to clients who face significant security and privacy challenges like censorship and network monitoring. However, we suggest that vendors can allow clients to *opt-in* to DoH services, and even offer clients with potentially useful information to help them decide (e.g., providing the user with data on how their web browsing performance would degrade if DoH was turned on). We note that vendors may already be rolling out DoH deployment in waves, for example, Firefox and Chrome on Android have turned on DoH by default for US clients [2, 36]. However, many vendors have not explicitly released their DoH rollout plans.

Improving DoH Resolution Services. Even when controlling for the resolution service used (e.g., Google, Cloudflare), the second largest factor in DoH slowdowns was the distance to the recursive resolver performing the resolution. We observe that different providers take significantly different approaches—Cloudflare, for example, has invested in significant geographic spread (146 PoPs) compared to Google, who has a relatively small number of PoPs (26) that handle significant geographic regions. One potential area of improvement for DoH performance may be to begin investing in small PoPs in areas with little development to reduce the time taken to get into the DoH provider network. However, as we show in Figures 5 and 6, having many PoPs is not enough on its own. For example, we observed significantly more PoPs for Quad9 than for any other provider in Sub-Saharan Africa, but clients in this region frequently use PoPs across the continent, or even across the world. In some cases, these PoP allocations may not be explicitly due to the resolution service itself, but rather that the service may rely on BGP anycast to perform routing, which has known inefficiencies [28]. Still, providers should ensure that clients are taking a full advantage of the PoPs nearby by continuing to improve their methods for assigning the optimal PoP to each client. Furthermore, while this seems likely to improve resolution times to some extent, nationwide bandwidth is *still* the largest factor that dictates DoH performance, and must be carefully considered before switching clients to DoH by default.

Cache Hits and Misses. Our study excluded the impact of caching when comparing DoH and Do53—the goal was to attribute the performance differences to transport protocols instead of domain names resolved. A drawback is that the results may not reflect each clients’ real-world performance (i.e., which involves both cache hits and cache misses). Rigorously comparing the performance under cache hit and cache miss is an interesting venue for future work. Intuitively, DoH is more “centralized” than Do53; it

would be interesting to study whether a more centralized cache implementation would lead to more or less cache hits, and how the caching performance eventually affect client experience.

Evaluating DoH Performance for Internet Applications. Previous studies of DoH performance have suggested that DNS is just a small part of web loading times and can even improve web page loading times on fast connections [21]. While this may be true for web browsing, DNS underpins almost all Internet communication—for example, software updates, instant messaging, and content delivery—and studying how DoH performance impacts other contexts remains an important area of future work.

Limitations. One core limitation of our study is the bias introduced by using a single proxy service, BrightData, for all of our measurements. We acknowledge this may introduce a bias towards users that are more technically savvy (e.g., ones using a proxy service in the first place). In addition, due to BrightData system restrictions, we could not study per-client differences for 11 countries, though several of them have been studied extensively in prior work [29]. Another limitation is the number of exit nodes available in a country varies in the BrightData network. In our study, we selected countries that had at least 10 unique clients. The uniqueness of the clients was ensured by the unique ID assigned by the Super Proxy. The number of clients per country varies from 10 to 282. This might skew our results due to some countries being underrepresented but our analysis show that the results are still statistically significant. Our study also only used a single authoritative name server in one location, whereas actual DNS performance depends on name servers located throughout the world. Our models in Section 6 did control for the distance to the name server, however, future work may want to vary name server location to simulate a more realistic DNS environment. Finally, our study only considers TLS 1.3, and clients that still use TLS 1.2 will have slower DoH performance overall. However, relative trends (e.g., between infrastructure investment and DoH performance) will likely remain consistent.

8 RELATED WORK

Measurements of DNS-over-Encryption. Our work follows from many measurement studies of encrypted DNS performance [8, 16, 18, 21, 22, 29, 33]. Our study complements these existing works by significantly increasing the coverage of vantage points (22,052 unique clients over 224 countries). This allows us to study the DoH performance around the world and examine correlated factors.

Lu et al. [29] conduct measurements in a large number of countries (100+). However, they cannot obtain the absolute DoH and Do53 resolution time with the clients’ default Do53 resolvers. Interestingly, they reported a reachability over 99% from exit nodes in China to Cloudflare and Quad9 in 2019. However, we observe that in 2021, 99% of the DoH queries sent from exit nodes in China were completely dropped. It is possible that related censorship policies in China have updated in the past two years.

With a focus on *DoT*, Doan et al. [16] obtains the absolute resolution times using 3.2K volunteer probes in the RIPE Atlas network. While they focus on a different encrypted DNS protocol (i.e., DoT), their study shares some similar observations with our study on DoH.

For example, they show that DoT generally has slower response times than Do53; when comparing different DoT resolvers, they also observe that Cloudflare and Google have better performance than Quad9. In contrast to [16], our study is of a larger scale (22K clients) and focuses on country-level analyses rather than *continent level* analyses. In doing so, they conclude that Cloudflare is the only resolver that exhibits consistent response times across continents, whereas we find that all resolvers (including Cloudflare) exhibit a high level of regional variance. We also further explore the potential reasons (e.g., economies, Internet infrastructures, PoPs) for cross-country differences, which are not studied in prior work.

Regarding other DNS-over-Encryption solutions, researchers have measured the adoption of DNSSEC and explored reasons behind the slow adoption rate [9, 10]. Their focus is adoption and (mis-)configurations rather than performance.

Security and Privacy of DNS-over-Encryption. DNS-over-Encryption provides certain security/privacy benefits but it is not necessarily resilient against all adversaries. Hoang et al. [19] find that, under encrypted DNS, the IP addresses (visible to ISP adversaries) may still reveal the websites that users visit. On a similar track, Siby et al. [44] demonstrate that DoH traffic can be fingerprinted to infer user activities. Huang et al. [24] show that encrypted DNS can be downgraded to plain text DNS by an adversary by exploiting the DoH implementation in browsers. As DNS-over-Encryption is on the verge to be widely adopted, such security and privacy risks should be carefully considered, and further research should be done to harden these solutions.

Disparities across Populations. We are not the first to consider how changes in protocols and network infrastructure impact different populations in varying ways. In 2010, Howard et al. [23] studied the so-called “digital divide” between the higher and lower income groups in the U.S. and Canada. They gathered empirical measurements demonstrating differences in the frequency and nature of Internet usage across income groups. More recently, Nielsen et al. [38] studied progress towards closing this divide, finding that there may be a tipping point around 50% Internet usage for many populations, after which Internet usage increases more rapidly. Quan et al. [41] studied which networks and regions experience the largest changes in devices online between day and night, finding that areas with lower per capita GDP experience a higher percent decrease in connected devices at night, possibly due to an increased emphasis on saving energy and money in those regions. These studies emphasize the importance in understanding how new technologies can affect different populations in substantially different ways.

9 CONCLUSION

The work studied the performance impact that a transition to DoH would have to residential clients around the globe. We devised a careful methodology, employing 22,052 clients across the world to collect our measurements. The resultant data paints a complex picture, with DoH providing performance benefits in certain regions and slowdowns in others. We then studied differences between four major public DoH providers, outlining differences in architecture and routing capabilities that may affect overall resolution performance. We also analyzed several explanatory variables correlated

with the performance impact of a switch from Do53 to DoH, finding that clients in countries with higher quality Internet infrastructure (faster speeds, more ASes) and clients from higher-income countries are less likely to experience a slowdown from a switch to DoH, and in many cases may experience a DoH *speedup*. This raises important questions about the asymmetrical effects of global DoH adoption, and should be studied further and weighed in DoH deployment decisions. We make our dataset available for further study, and we hope our findings will help inform the DoH community as adoption continues to accelerate.

ACKNOWLEDGMENTS

We thank our shepherd Aaron Schulman and the anonymous reviewers for their constructive comments and suggestions. This work was supported in part by the National Science Foundation (NSF) under grant #2030859 to the Computing Research Association for the CIFellows Project, and grant #2030521. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agencies.

REFERENCES

- [1] Apple. 2020. Enable encrypted DNS. <https://developer.apple.com/videos/play/wwdc2020/10047/>.
- [2] Kenji Baheux. 2020. A safer and more private browsing experience on Android with Secure DNS. <https://blog.chromium.org/2020/09/a-safer-and-more-private-browsing.html>.
- [3] World Bank. 2021. <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>.
- [4] BrightData. 2021. Bright Data (formerly Luminati Network). <https://brightdata.com/>.
- [5] BrightData. 2021. Bright Data’s super proxy servers. <https://brightdata.com/proxy-types/super-proxy>.
- [6] Martin Brinkmann. 2020. Chrome 83: rollout of DNS over HTTPS (Secure DNS) begins. <https://www.ghacks.net/2020/05/20/chrome-83-rollout-of-dns-over-https-secure-dns-begins/>.
- [7] Content by Rodney. 2021. How to Enable Encrypted DNS on iPhone iOS 14. <https://rodneylab.com/how-to-enable-encrypted-dns-on-iphone-ios-14/>.
- [8] Timm Böttger, Felix Cuadrado, Gianni Antichi, Eder Leão Fernandes, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2019. An Empirical Study of the Cost of DNS-over-HTTPS. In *ACM Internet Measurement Conference*.
- [9] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. 2017. A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In *USENIX Security Symposium*.
- [10] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. 2017. Understanding the Role of Registrars in DNSSEC Deployment. In *ACM Internet Measurement Conference*.
- [11] Google Cloud. 2021. <https://cloud.google.com/cdn/docs/locations>.
- [12] Cloudflare. 2021. Android Pie and later supports DNS over TLS. <https://developers.cloudflare.com/1.1.1.1/setup-1.1.1.1/android>.
- [13] Cloudflare. 2021. *Cloudflare 1.1.1.1*. <https://1.1.1.1>
- [14] Federal Trade Commission. 2021. <https://www.fcc.gov/consumers/guides/broadband-speed-guide>.
- [15] Casey Deccio and Jacob Davis. 2019. DNS Privacy in Practice and Preparation. In *ACM International Conference on emerging Networking EXperiments and Technologies*.
- [16] Trinh Viet Doan, Irina Tsareva, and Vaibhav Bajpai. 2021. Measuring DNS over TLS from the Edge: Adoption, Reliability, and Response Times. In *Passive and Active Measurement Conference*.
- [17] Google. 2019. *Google Public DNS*. <https://developers.google.com/speed/public-dns>
- [18] Brian Haberman and Catherine Master. 2017. DNS-over-TLS Measurements with RIPE Atlas Probes. <https://datatracker.ietf.org/meeting/102/materials/slides-102-dprive-dns-over-tls-measurements-with-ripe-atlas-probes-01>.
- [19] Nguyen Phong Hoang, Arian Akhavan Niaki, Nikita Borisov, Phillipa Gill, and Michalis Polychronakis. 2020. Assessing the Privacy Benefits of Domain Name Encryption. In *ACM ASIA Conference on Computer and Communications Security*.
- [20] Austin Hounsel, Kevin Borgolte, Paul Schmitt, Jordan Holland, and Nick Feamster. 2019. Analyzing the Costs (and Benefits) of DNS, DoT, and DoH for the Modern

- Web. In *Applied Networking Research Workshop*.
- [21] Austin Hounsel, Kevin Borgolte, Paul Schmitt, Jordan Holland, and Nick Feamster. 2020. Comparing the Effects of DNS, DoT, and DoH on Web Performance. In *The ACM Web Conference*.
- [22] Austin Hounsel, Paul Schmitt, Kevin Borgolte, and Nick Feamster. 2021. Can Encrypted DNS Be Fast?. In *Passive and Active Measurement Conference*.
- [23] Philip N Howard, Laura Busch, and Penelope Sheets. 2010. Comparing digital divides: Internet access and social inequality in Canada and the United States. *Canadian Journal of Communication* 35, 1 (2010).
- [24] Qing Huang, Deliang Chang, and Zhou Li. 2020. A Comprehensive Study of DNS-over-HTTPS Downgrade Attack. In *USENIX Workshop on Free and Open Communications on the Internet*.
- [25] Geoff Huston. 2018. DOH! DNS over HTTPS explained. <https://blog.apnic.net/2018/10/12/doh-dns-over-https-explained/>.
- [26] IPInfo. 2021. <https://ipinfo.io/countries>.
- [27] Bind9 ISC. 2021. *Bind9 Name Server - ISC*. <https://www.isc.org/bind/>
- [28] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet anycast: performance, problems, & potential. In *ACM Special Interest Group on Data Communication*.
- [29] Chaoyi Lu, Baojun Liu, Zhou Li, Shuang Hao, Hai-Xin Duan, Mingming Zhang, Chunying Leng, Ying Liu, Zaifeng Zhang, and Jianping Wu. 2019. An End-to-End, Large-Scale Measurement of DNS-over-Encryption: How Far Have We Come?. In *ACM Internet Measurement Conference*.
- [30] Luminati. 2021. *Monetize your active and inactive users by becoming a Bright SDK partner*. <https://luminati.io/sdk>
- [31] Mauro Huc. 2021. How to enable DNS over HTTPS (DoH) on Windows 11. <https://pureinfotech.com/enable-dns-over-https-windows-11/>.
- [32] Maxmind. 2021. <https://www.maxmind.com>.
- [33] Patrick McManus. 2018. Firefox Nightly Secure DNS Experimental Results. <https://blog.nightly.mozilla.org/2018/08/28/firefox-nightly-secure-dns-experimental-results/>.
- [34] Microsoft. 2019. <https://techcommunity.microsoft.com/t5/networking-blog/windows-will-improve-user-privacy-with-dns-over-https/ba-p/1014229>.
- [35] P. Mockapetris. 1987. Domain names - implementation and specification. RFC 1035. <https://doi.org/10.17487/RFC1035>
- [36] Mozilla. 2020. <https://blog.mozilla.org/blog/2020/02/25/firefox-continues-push-to-bring-dns-over-https-by-default-for-us-users/>.
- [37] NextDNS. 2021. *NextDNS*. <https://nextdns.io/>
- [38] Morten Meyerhoff Nielsen, Ibrahim Kholilul Rohman, and Nuno Vasco Lopes. 2018. Empirical Analysis of the Current Digital Divides since 2010. In *International Conference on Theory and Practice of Electronic Governance*.
- [39] Ookla. 2021. <https://www.speedtest.net/global-index>.
- [40] Quad9. 2018. *DoH with Quad9 DNS Servers*. <https://www.quad9.net/news/blog/doh-with-quad9-dns-servers/>
- [41] Lin Quan, John Heidemann, and Yuri Pradkin. 2014. When the Internet Sleeps: Correlating Diurnal Networks with External Factors. In *ACM Internet Measurement Conference*.
- [42] Eric Rescorla. 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446. <https://doi.org/10.17487/RFC8446>
- [43] RIPE NCC. 2021. *What is RIPE Atlas?* RIPE NCC. <https://atlas.ripe.net/about/>.
- [44] Sandra Siby, Marc Juarez, Claudia Diaz, Narseo Vallina-Rodriguez, and Carmela Troncoso. 2019. Encrypted DNS → Privacy? A Traffic Analysis Perspective. In *ISOC Network and Distributed Systems Security Conference*.
- [45] Systemd. 2021. Add support for DNS-over-HTTPS to systemd-resolved. <https://github.com/systemd/systemd/issues/8639>.
- [46] TrendMicro. 2018. Shining a Light on the Risks of HolaVPN and Luminati. <https://www.trendmicro.com/vinfo/hk-en/security/news/cybercrime-and-digital-threats/shining-a-light-on-the-risks-of-holavpn-and-luminati>.
- [47] Liang Zhu, Zi Hu, John S. Heidemann, Duane Wessels, Allison Mankin, and Nikita Somaiya. 2015. Connection-Oriented DNS to Improve Privacy and Security. In *IEEE Security and Privacy Symposium*.

A ETHICS

Our measurement methodology involves the use of the BrightData proxy network and RIPEAtlas measurement platform. Below, we discuss the key ethical aspects related to our experiments. First, BrightData is a commercial platform. We purchased their services and strictly followed their Terms of Service when running our experiments. We also had significant conversations with the BrightData team about our experiments ahead of time and received explicit approval from them that we could conduct our measurements on their platform. Second, the exit nodes of the BrightData network are

Resolver	Metric	Coef. (ms)	Scaled Coef. (ms)
Cloudflare	GDP	2e-4*	4.14*
	Bandwidth	-1.4	-85.3
	Num ASes	-6.3e-2	-85.8
	Nameserver Dist.	1.23e-2	32.7
	Resolver Dist.	9.33e-2	155.7
Google	GDP	-5.18e-5*	-1.07*
	Bandwidth	-0.95	-56.8
	Num ASes	-5.12e-2	-69.7
	Nameserver Dist.	1.54e-2	40.87
	Resolver Dist.	8.48e-2	140.02
NextDNS	GDP	-9.66e-4*	-19.9*
	Bandwidth	-2.32	-138.3
	Num ASes	-7.34e-2	-99.8
	Nameserver Dist.	6.48e-3	17.2
	Resolver Dist.	6.78e-2	111.99
Quad9	GDP	-1.05e-3*	-21.6*
	Bandwidth	-2.1	-124.1
	Num ASes	-3.6e-2	-49.1
	Nameserver Dist.	1.1e-2	27.8
	Resolver Dist.	3.4e-2	56.0

Table 6: Linear Modeling of DNS Performance by Resolver—

We show the results of our linear modeling split by each resolver for the delta between a single DoH request compared to Do53. All results are statistically significant with $p < 0.001$ unless indicated by an asterisk (*).



Figure 8: Clients in Our Dataset—A map showing the clients we used to conduct our Do53 and DoH measurements. Clients in our dataset span 22,052 unique IP addresses across 224 total countries and territories.

recruited/enrolled by the platform. The exit nodes and the BrightData platform have agreements to route traffic through the exit nodes (in exchange for free VPN services). Third, our experiments only involve generating DNS queries to benign DNS resolvers to query benign domain names (under our control). This experiment does not introduce any harm to the proxy service or the exit nodes. We note that we do not ever store raw client IP addresses in our study. Any geolocation lookups presented in this work are based on the /24 of the IP address. We do log the IP addresses of the public recursive resolvers used to perform DoH queries, however, we take careful note not to inspect any potentially sensitive client data (e.g., client IPs present in the ECS-client-subnet DNS extension).




Figure 9: Per-Client Distance to Servicing DoH PoP—For each resolver, we present the distance for each client to the PoP that was used to perform the resolution. Providers excel in varying regions. Quad9, for example, assigns PoPs which are relatively close for clients in the southern part of Africa, but seems to underperform other providers by this metric in southern Brazil and Argentina.

Finally, our use of RIPE Atlas is in line with their terms of service and involves just minimal testing for ground truth validation experiments.

B DATASET

We provide some additional figures characterizing our dataset. Figure 8 shows a map with all of the Maxmind-located clients we used in this study (locating them based on their /24). Figure 9 shows

these clients by resolver, coloring them based on the geographic distance to the DoH resolver point-of-presence they used.

C REGRESSION ANALYSIS

In addition to conducting aggregate regressions, we also compute linear regressions filtered by each DoH resolver provider. We show the results in Table 6.