SAC109:
The Implications of DNS over HTTPS and DNS over TLS

A report from the ICANN Security and Stability Advisory Committee (SSAC)
12 March 2020

Preface

This is a report of the ICANN Security and Stability Advisory Committee (SSAC). The SSAC focuses on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., pertaining to the correct and reliable operation of the root zone publication system), administrative matters (e.g., pertaining to address allocation and Internet number assignment), and registration matters (e.g., pertaining to registry and registrar services). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no authority to regulate, enforce, or adjudicate. Those functions belong to other parties, and the advice offered here should be evaluated on its merits.

Table of Contents

Executive Summary

Encrypted DNS technologies, including DNS-over-HTTPS (DoH) and DNS-over-TLS (DoT), are recent protocols developed for the primary purpose of enhancing user privacy. They accomplish this in several ways, including encrypting their traffic in transit and permitting DNS resolver selection and resolution in applications.

Major browser vendors, Internet Service Providers (ISPs), and others are deploying support for these technologies. Their deployment brings a number of possible implications, both positive and negative, to the ICANN community, operators and users of the DNS, and Internet users.

This report analyzes the initial effects of these technologies by identifying some groups whose online experiences around privacy could change with the deployment of these technologies. Detailed analysis of effects will have to wait for more widespread deployment and measurement. This report discusses implications occurring now, and raises some longer-term questions for the future. This report frames the issues from the perspectives of interested parties, with the understanding that the issues are nuanced, and that readers coming from different perspectives will have different sensitivities: readers from two different perspectives are likely to view a single issue in two different ways.

The intended audience for this report is both the ICANN community and the greater Internet community. This includes network operators, DNS software implementers, policy makers, and concerned Internet users.

Note to the Reader

The SSAC began its work on this document with a good deal of discussion that culminated with some important realizations.

The issues raised by DoH and DoT and by the implementation and deployment choices that exist for them are not straightforward. It is not possible to lay them out with universally agreed-upon "right" and "wrong" labels. We realized that we were not going to get consensus on clear, strong statements such as, "More privacy is always better," or "More encryption is always better." Nor could we make clear, strong statements about trust models or the like that we would all agree with, because we all look at it in different ways.

Instead, the SSAC decided the best way to approach this would be to describe things from different points of view, which we've called perspectives, showing how different perspectives view these issues differently. The hope is that readers can see the different sides of each issue and understand why not everyone is making the same choices about all this.

We focused on explaining things for the community we primarily serve: the ICANN community. That meant we said things that are also said elsewhere, because we could not expect others to have the background that we do, nor have done the research that we have. Some of this is said in various articles on the Internet, in IETF Internet Drafts, and in other venues. What we say in this document should support those and should be supported by those, but isn't made unnecessary by those.

1 Introduction

This report is intended to help the ICANN community understand the implications of a relatively new set of protocols that transport DNS queries and responses over encrypted channels. The report will explain the technical changes introduced with these protocols, and will explore in-depth the security and stability effects of how these changes affect different actors and the DNS as a whole. The target audience for this work is the ICANN community and the larger Internet community.

The protocols specifically covered here are DNS over HTTPS (DoH) and DNS over TLS (DoT). Specifications for DNS over other new transport protocols — such as DNS over QUIC[1] and DNS over DTLS[2] — exist, are under development, or are expected to appear. Most of what will be discussed herein will also be applicable to other such transports, but will not be directly covered in this publication.

The DNS is changing all the time, but some changes in how the DNS operates have far-reaching and subtle implications. There are several steps involved in the process of creating, publishing, and retrieving DNS data. While the ultimate purpose of publishing and retrieving data does not change, the protocols and configuration conventions used in the delivery of that data can change quite a bit. Each part of the DNS may be operated by a different entity, and changes in how the players and the parts fit together can have not only operational implications that change bits on the wire, but policy and economic effects as well.

In the traditional model of DNS resolution, a DNS library is included in operating systems. While the resolver used by this library is sometimes configured by the end user, it is more often configured by the service provider through the use of the Dynamic Host Configuration Protocol (DHCP). The configured resolver is mostly a system-wide setting and generally not application-specific.

There have been recent disruptions of the traditional model on multiple fronts:

- New technical standards and implementations have been developed to convey DNS queries and responses over alternative transport protocols, examples of which are HTTPS,[3] TLS,[4] DTLS,[5] and QUIC.[6] Such efforts move away from unencrypted UDP and

[1] See Huitema, C., Shore, M., Mankin, A., Dickinson, S., Iyengar, J., "Specification of DNS over Dedicated QUIC Connections", draft-huitema-quic-dnsoquic-07, September 2019, https://datatracker.ietf.org/doc/draft-huitema-quic-dnsoquic/

[2] See RFC 8094

[3] See RFC 8484

[4] See RFC 7858

TCP (the traditional protocols used to transport DNS traffic) and appear to be driven by privacy, confidentiality, security, and robustness considerations, as well as a desire for increased levels of control over the retrieval of DNS information for client applications.

- At the time of publication, open public resolver services not operated by ISPs handle roughly 16% of all DNS resolution on the Internet.[7] Examples of open public resolver operators include; Google,[8] OpenDNS[9] and 114DNS.[10] Such efforts consolidate DNS-based user behavior and consequently introduce new and different privacy questions related to user data collection.
- Vendors of browsers and other applications have incentives to embed addresses of resolvers directly into their applications, thereby bypassing the traditional model of using the resolver(s) configured in the operating system, and instead creating application specific resolution behaviors. An example of this is Mozilla Firefox, which at the time of publication, is sometimes pre-configured with a DoH capable resolver compliant with Mozilla's Trusted Recursive Resolver Program[11] (TRR).
- Mechanisms that use the DNS as a control point are continuing to be developed and implemented. These include DNS-aware firewalls and monitoring tools, as well as some forms of IPv6 transition technologies. These may be local to end-users or centralised, and are often deployed at the direct request of the end-user or to enforce a policy of the local network.

The confluence of these technologies is fundamentally changing some aspects of DNS resolution. Study is needed on the long term technical, operational, political, and economic implications for the DNS industry and the associated policy implications for the ICANN community. Design changes to the Internet and the DNS often happen through uncoordinated action by a diverse group of actors with differing incentives.[12] This makes it difficult to predict what outcomes will result from the development of new technologies and their deployments.

[5] See RFC 8310

[6] See Huitema, C., Shore, M., Mankin, A., Dickinson, S., Iyengar, J., "Specification of DNS over Dedicated QUIC Connections", draft-huitema-quic-dnsoquic-07, September 2019, https://datatracker.ietf.org/doc/draft-huitema-quic-dnsoquic/

[7] See APNIC, Use of DNS Resolvers for World, https://stats.labs.apnic.net/rvrs/XA?hc=XA&hl=1&hs=1&ht=0&w=30&t=0&s=0

[8] See Google Public DNS, https://developers.google.com/speed/public-dns/

[9] See Cisco OpenDNS, https://www.opendns.com/

[10] See 114DNS, https://www.114dns.com/

[11] See Mozilla's Trusted Resolver Program, https://wiki.mozilla.org/Trusted_Recursive_Resolver

[12] See Clark, D., Wroclawski, J., Sollins, K., Braden, R., "Tussle in Cyberspace: Defining Tomorrow's Internet", IEEE/ACM Transactions on Networking, VOL. 13, NO. 3, JUNE 2005, https://dl.acm.org/doi/10.1109/TNET.2005.850224

1.1 The Internet Engineering Task Force

The Internet Engineering Task Force (IETF) is the standards development organization responsible for many of the standards that make the Internet work and keep it running, and its responsibilities include the standards for DNS. Part of the IETF's remit is to make sure that its standards support integrity and security of the Internet. The IETF community has been increasingly concerned, over the years, with the security and privacy aspects of the Internet protocols.

The revelations of Edward Snowden in 2013 regarding pervasive government surveillance of the global Internet infrastructure provoked a much deeper awareness about the vast amount of data every user was making available to anyone with the technical ability to monitor their interactions with the network. This information could be used for any purpose the watcher desired, unknown to the user, and the monitoring could not be disabled. Both RFC 7258[13] and RFC 7624[14] resulted from the discussions of these issues, and those documents explain the IETF community's views in some depth.

Since the publication of those documents, the IETF has further increased its focus on security, privacy, and confidentiality, resulting in more built-in encryption along with stronger recommendations for the use of encryption where it had been optional. In this regard, the IETF looks to adhere to a rough set of principles that it has set for Internet protocols, trying to design for an open yet secure Internet in which parties can be confident. Among those principles is respect for the privacy of user activities: that user interactions with the network should be private by default, even though many of the basic protocols, including DNS, were written before this issue had attracted the interest it has today. While the underlying infrastructure of the Internet can't prevent users from giving up their privacy, by choice or duress, it can at least not require them to do so as a condition of participating in the network.

That said, the IETF does not, itself, create technology nor decide what technology will be deployed in a given application or network. Application builders and network operators do that, and they make their own choices as to what to support and how it will be configured. It is often noted that there is no "Internet police", and what the IETF provides are building blocks and recommendations.

DNS Security Extensions (DNSSEC) were first standardized in the late 1990s, and provide integrity protection for responses to DNS queries. When DNSSEC is in use, applications can be

[13] See RFC 7258
[14] See RFC 7624

sure that the content of query responses are not fraudulent and have not been altered. However, DNSSEC provides no confidentiality protection. DNS responses themselves contain only public information, but the specific queries that are sent, and by whom, can leak enormous amounts of potentially damaging user data. Once monitoring and data exfiltration were seen as important threats both within and outside of the information security community, modernizing the DNS protocol to provide better confidentiality for users became a priority.

The IETF DNS PRIVate Exchange (DPRIVE)[15] working group hosted most of the work of DNS protocol hardening to support confidentiality, including DoT, with the DNS Over HTTPS (DoH) [16] working group later set up specifically for work on the DoH specification. DoH and DoT transactions are encrypted, providing confidentiality of DNS queries and responses in transit on the network. DoH also provides some ability to hide DNS transactions in a stream of activity related to other protocols and content, defeating some metadata collection. They do nothing to protect information at the endpoints of a transaction, but they can protect it from interception in between.

Recently "Birds of a Feather" sessions at IETF 105[17] and 106[18] reviewed the current state and discussed what additional work might be needed to make DoH and DoT deployable, addressing both protocol extensions and "best practices" guidance to developers and operators. Work has also begun on an Internet Draft focusing on issues and risks related to the centralization of DoH deployments.[19] Work in this area remains ongoing.

2 Overview of this document

The remainder of this document will cover the following:

- A comparison of the transport protocols intended to improve DNS privacy, focusing on the standardization and deployment status of specific technologies.
- The direct and indirect effects of these technologies on several different groups of stakeholders. We describe the effects from the viewpoint of each group without favoring the outlook of any one group.

[15] See DNS PRIVate Exchange (dprive) Working Group, https://datatracker.ietf.org/wg/dprive/about/

[16] See DNS Over HTTPS (doh) Working Group, https://datatracker.ietf.org/wg/doh/about/

[17] See Applications Doing DNS (add) Proposed Working Group, https://datatracker.ietf.org/wg/add/about/

[18] See Application Behavior Considering DNS (ABCD) BoF Meeting Minutes, 19 November 2019, https://datatracker.ietf.org/meeting/106/materials/minutes-106-abcd-00

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

- The issue of who decides which resolvers are used by hosts and specific applications, and what implications arise from these decisions.
- Potential implications on the namespace due to DNS stub resolution moving to applications.
- Conclusions related to the deployment and adoption of these technologies for users and policy makers.

DNSSEC is out of scope for this document; it provides an important, but separate, set of protections. DNSSEC provides integrity protection, not confidentiality. Both are important and can be deployed in parallel, but they solve different problems.

Additionally, this document will not take a position on the merits of DNS blocking, a topic the SSAC has previously addressed.[20,21] Also out of scope are technologies that can be used to improve user privacy in normal DNS interactions, for example Query Name Minimisation.[22] This document focuses on transport-layer technologies, not packet contents. Finally, efforts to standardize and deploy encrypted transport between recursive and authoritative DNS servers are out of scope for this document.

3 Comparison of DNS over HTTPS and DNS over TLS

DNS queries and responses are traditionally transported in clear text (unencrypted) over the underlying UDP or TCP protocol.[23] DoT[24] specifies that DNS queries and responses be transported over TLS, while DoH[25] relies on HTTPS, the protocol used for secure (i.e., encrypted) web pages.

It's important to understand that the protocols DoT and DoH change only the underlying transport channel for the DNS queries and responses. They do not themselves make any changes to the queries and responses (the DNS protocol), nor to the mechanism used to build the responses (DNS resolution).[26] Any changes beyond the use of encrypted transport are a result of implementation and deployment choices, and not of the DoT and DoH protocols. The use of encryption provides a level of confidentiality for the DNS queries and responses, in that someone

[20] See SAC050
[21] See SAC056
[22] See RFC 7816
[23] See RFC 1035
[24] See RFC 7858
[25] See RFC 8484
[26] A minor exception to this is that RFC 8484 section 4.1 states that DoH clients SHOULD always use a DNS message ID of 0 to facilitate easier HTTPS caching.

who can surveil Internet traffic will only see encrypted traffic and who sent it, but will not be able to determine what the content of the queries and responses are.

3.1 Comparison of Possible Deployments

The four figures below represent four possible deployment models for traditional DNS, DoT, and DoH. They are representative of the deployment models discussed during protocol development, and that have been seen in early deployments, but are not meant to enumerate all possible deployments of these technologies.

In each figure, the application and user are shown on the left and authoritative DNS servers are shown on the right. The arrows trace DNS queries from the application to authoritative servers. Green dotted lines represent unencrypted transport, while red solid lines represent encrypted transport. Where encrypted transport is used, intermediaries who cannot decipher the query are shown as being traversed instead of taking an active role in the resolution path.

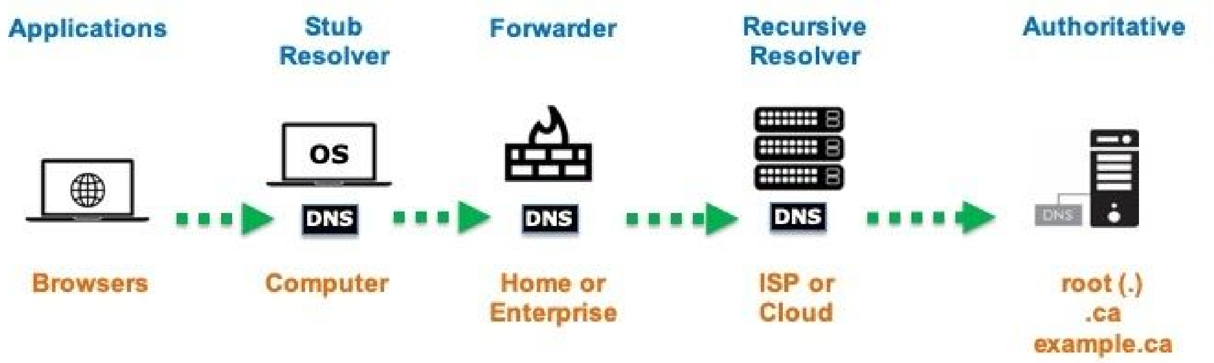

Figure 1: Possible Traditional DNS Deployment

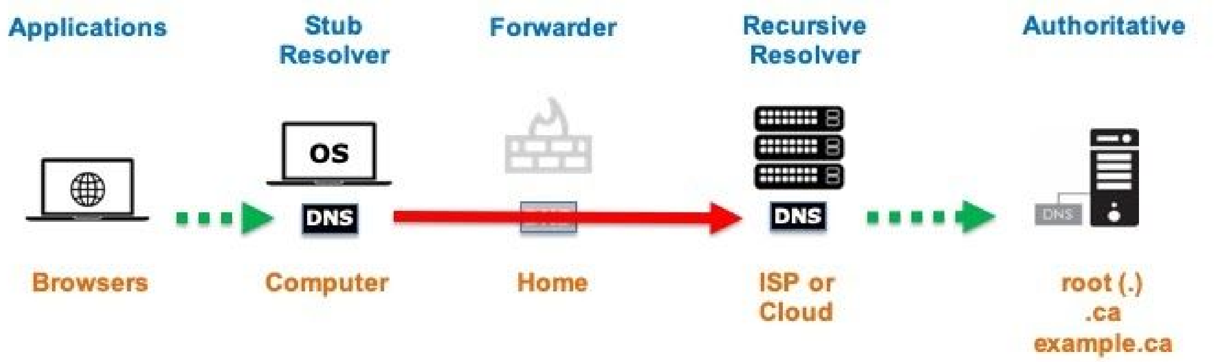

Figure 2: Possible DNS over TLS Deployment in a Home Network

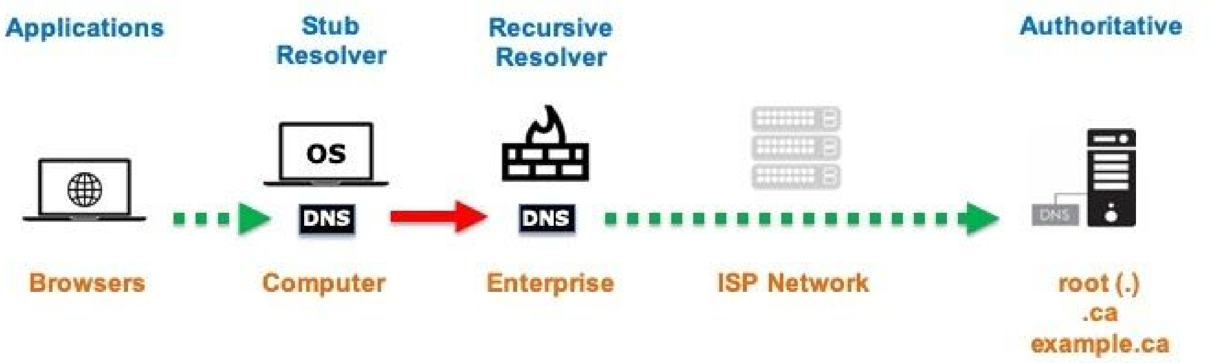

Figure 3: Possible DNS over TLS Deployment in an Enterprise Network

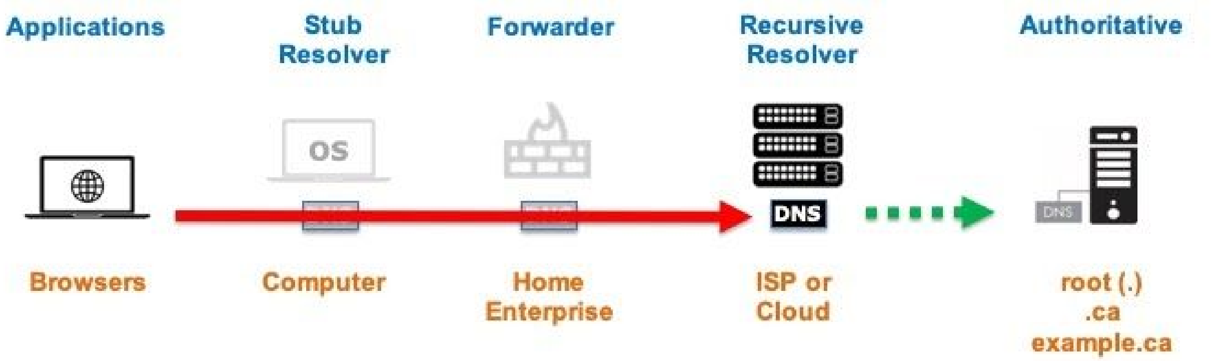

Figure 4: Possible DNS over HTTPS Deployment

3.2 Differences Between DoH and DoT

Both DoH and DoT make use of the same encryption methods for queries and responses. However, DoT will likely be implemented in operating systems, whereas DoH will likely be implemented in both applications and operating systems. Users who enable DoT in their operating system will likely start using it through the explicit change of defaults.

With DoH this will likely not be the case. Many end-users of DoH will start using it by default, perhaps unknowingly, as the applications they regularly use start to implement it. Many applications, including web browsers, already use HTTPS as their primary mode of transport. Thus, it will be easier for them to implement DoH than DoT in their applications, and they can be more certain that deploying DoH will cause fewer connectivity problems for their user base.

This defacto segregation between applications tending to support DoH and platform service libraries tending to support DoT appears to reflect a common informal position that DoH may be more readily integrated into a richer context of an HTML session that can be used to control an application's use of the DNS name resolution function. On the other hand, DoT is seen as a simple addition to the TCP transport of the DNS protocol that adds a secure association between the platform's stub resolver and its configured recursive resolver.

Another important distinction between DoT and DoH is that DoT uses well known TCP port 853 by default, while DoH makes use of HTTPS port 443. Thus, DoT can be identified and blocked by intermediaries such as network administrators. Many firewalls block all ports by default and only whitelist specific well-known ports such as port 443 for HTTPS. DoT adoption may face obstacles because deployed firewalls and other middle boxes will not permit traffic on TCP port 853 to leave the network, thereby blocking the query traffic.

DoH is designed to be indistinguishable from normal HTTPS to the same server, making it harder to block by intermediaries. All DoH traffic cannot be blocked without potentially blocking other important HTTPS traffic if a server handles both DoH and regular web traffic. Research in this area is ongoing, with recent studies showing that through traffic analysis encrypted DNS traffic can sometimes be distinguished from other HTTPS traffic, that sometimes it is possible to determine queried domain names, and even that sometimes it is possible to determine which websites users are visiting.[27,28,29] These techniques will likely only get better with time.

The encryption of traffic in transit removes a particular set of passive attacks and monitoring techniques. However, participating resolvers themselves will still have full access to the queries and responses passing through them and are equally useful for surveillance and filtering activities. One concern with the use of end-point to resolver communications is that there may be less anonymization than the enterprise or ISP resolvers may have provided using traditional

[27] See Traffic Analysis still possible when using DoT or DoH, https://blog.apnic.net/2020/01/30/traffic-analysis-still-possible-when-using-dot-and-doh/

[28] See Houser, R., Li, Z., Cotton, C., Wang, H., "An investigation on information leakage of DNS over TLS", CoNEXT '19: Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, December 2019, 123–137, https://doi.org/10.1145/3359989.3365429

[29] See Siby, S., Juarez, M., Diaz, C., Vallina-Rodriguez, N., Troncoso, C., "Encrypted DNS ⇒ Privacy? A Traffic Analysis Perspective", October 2019, https://arxiv.org/pdf/1906.09682.pdf

DNS.[30] Particularly with DoH, depending on how it is deployed, the same operator may have DNS, HTTPS, and other metadata to relate DNS queries to users and their activities.

3.3 Dependencies on Traditional DNS

For the time being, with both DoH and DoT, resolution depends on traditional DNS over UDP or TCP port 53 to function on the connection between a recursive and an authoritative resolver, since neither protocol is standardized for authoritative DNS servers. Additionally, DoH is dependent on HTTPS, which is itself dependent on DNS for resolving the host part of Uniform Resource Identifiers (URIs).[31] While it is possible to use IP addresses in URIs instead of domain names, URIs for DoH services will most likely use domain names. Domain names in URIs can be more easily authenticated than IP addresses due to restrictions placed on X.509 certificate issuance, and using domain names instead of IP addresses permits the DoH operator to utilize virtual hosting. DoH clients will most likely use traditional DNS to resolve domain names in URIs and will use the web's Public Key Infrastructure (PKI) to authenticate the servers they contact.

Like DoH, DoT clients requiring strict privacy must authenticate the servers they contact. Unlike DoH, none of the authentication mechanisms for DoT depend on some other method to perform DNS resolution (i.e., traditional DNS).[32] All authentication mechanisms available for DoT allow for either the bootstrapping of authentication using DNS data provided over DoT, or preconfigured information in the DoT client.

4 Perspectives on DNS Encrypted Transport in Applications

This section provides different perspectives on the use of DoH and DoT by describing how the protocols and their deployment may affect different groups of people. We have chosen these specific groups to illustrate the divergent opinions on DoH and DoT and to highlight that there is no simple answer to many of the policy questions around these protocols and their deployment. For each group, we discuss the most relevant implications of DoT and DoH for that group, what about DoH and DoT they may see as positive or useful, what concerns they may have, and how it may alter the status quo.

[30] Traditional DNS resolvers hosted by enterprises and ISPs receive queries from many different hosts, most of which will be answered from the resolver's cache. Therefore queries sent from these resolvers cannot easily be associated with the queries they receive from hosts. This inability to associate inbound queries with outbound queries provides a basic level of anonymity for hosts using the resolver.

[31] See RFC 3986

[32] See RFC 8310

Some of the groups listed in this section filter or block DNS traffic for different reasons, and deployments of DoH and DoT will affect them in different ways. For example, parents (section 4.1) may block traffic to protect their children, and enterprise network managers (section 4.2) may block traffic to comply with local laws or regulations, or prevent the exfiltration of proprietary data. Different users within each group will experience the implications of new DNS technologies differently. The categories of user groups listed in this section are not meant to be exhaustive, nor are they able to capture the implications of new DNS technologies on all users in each group.

Not all computer users have the same level of knowledge nor are they equally comfortable with changing default settings on their computers, and some applications that set their own DNS behavior do not even allow users to change default values. Default DNS settings will affect different classes of users in different ways. Some users will be comfortable modifying their DNS settings, while others will not be. Many of the mechanisms deployed today to filter DNS responses can be bypassed by determined users. Thus, it is worth noting that DNS filtering in use today can be bypassed, especially by users that have “admin” access to their workstations. For example, a user could install a Virtual Private Network (VPN) browser extension, use an alternative browser which supports Tor,[33] or simply install a VPN application. This is especially true for Bring Your Own Device (BYOD) scenarios, or when users have devices (i.e., tablets, smartphones) where turning off WiFi causes these devices to use the cellular provided DNS service.

For more detailed information on DNS blocking, how it works, and how it affects users of the DNS, please refer to the SSAC's previous reports on the subject.[34,35]

4.1 Parents

For almost as long as residential Internet service has existed, services designed to control children’s access to the Internet have been marketed towards parents, guardians, and educational institutions. These services are offered by third parties and sometimes directly by residential ISPs.[36,37,38,39]

[33] See The Tor Project, https://www.torproject.org/

[34] See SAC050

[35] See SAC056

[36] See Kapersky Safe Kids, https://usa.kaspersky.com/safe-kids

[37] See Cisco OpenDNS DNS Security Services, https://www.opendns.com/home-internet-security/

[38] See BT Parental Controls: Keep your children safe online, https://home.bt.com/tech-gadgets/internet/broadband/stay-safe-with-bt-parental-controls-11363887238413

[39] See Internetmatters.org Parental Controls, https://www.internetmatters.org/parental-controls/

Opinions on whether, and how much, to control a child's access to the Internet vary. However, it is uncontroversial to observe that objectionable material exists on the Internet, even without a single definition of what, specifically, is objectionable. It’s similarly straightforward to observe that parents may routinely mediate the content available to their children while they are young, regardless of whether that content is accessible by other means.

While it is technically possible to block access to particular sources of content by blocking access to known IP addresses, IP addresses change. For example, a publisher of particular content might choose a different hosting provider, or a content delivery network might make content available on different parts of their infrastructure, thereby requiring new IP addresses for the servers hosting the content. Since tracking these changes is difficult and impractical at scale, and since content is most often located using URIs with domain names, managing access based on domain names is far easier and, perhaps unsurprisingly, a far more common mechanism.

Since the set of domain names that map to objectionable content vary, the filtering required to apply a policy that blocks objectionable content must be dynamic. A parent may want to apply a filtering policy across many devices and many different applications installed on those devices. The easiest way to do this is with the DNS, because if users cannot find the IP addresses of the servers hosting the content, they cannot access the content.

The router or routers that connect a home network to the Internet (“home gateway”) provide a possible point of configuration and control for such a policy to be applied. However, the ability to inspect and modify DNS responses in order to apply a content policy depends on that home gateway’s ability to observe unencrypted (cleartext) DNS traffic. There are two ways for this to happen:

1. The DNS resolver service used by devices and applications within the home network is provided by the home gateway. Devices and applications might make use of any available transport protocol, including DoH and DoT.
2. Devices and applications use a DNS resolver service that is located elsewhere, such as a centralised service reached over the Internet. In this case the home gateway only has the opportunity to mediate DNS traffic using transport that is unencrypted, or that is encrypted in a way that explicitly allows it to be intercepted.

Case 2 is problematic for a home network, as described and shown in Figure 2. While some types of encrypted transport might simply be blocked without significant blocking of other traffic (as with DoT), blocking DoH is not generally practical since it is more difficult to discern DoH from other HTTPS traffic that comprises the vast majority of non-objectionable content.

Blocking DoH traffic by blocking IP addresses of well known DoH resolvers is possible for centralised DoH resolvers whose addresses are well known, but this is not a universal solution.

The widespread deployment of DoH, not served by a resolver within the domain of control of the parent or ISP, will reduce or eliminate the ability for parents to apply content filtering policy using these kinds of services in use today, which rely upon DNS filtering.

4.2 Enterprise Network Managers

Enterprise networks can include many different kinds of organizations found around the world including; corporations, municipalities, university campuses, hospitals, and military bases. What usually defines an enterprise network is that, as a network, it does not provide a generalized transport network for the Internet. Internet traffic is either generated in the enterprise network or it is terminated in the enterprise network. It does not transit through the enterprise network. The primary purpose of enterprise networks is to provide access for their in-house users, typically employees.

Many enterprises have a obligation to understand and control the traffic on their networks. This obligation could be limited to owned assets or could be broadened to include any device that any user brings onto the enterprise's network. Failure to do so can expose them to a variety of problems, ranging from loss (the exfiltration of data), to regulatory penalties (certain industries are required to monitor their network traffic), to HR incidents (employees visiting sites that are inappropriate or impermissible for the workplace).

A common mechanism employed by enterprises to control the traffic on their networks is through the management of DNS. If a domain name points to a resource that is forbidden by the company’s IT policy, then the resource can be blocked (or at least detected) through DNS management, without having to inspect all web traffic in more detail. Many organizations have implemented DNS firewalls that utilize threat intelligence feeds and their own research to block domains known to be hosting and/or controlling phishing, malware, data exfiltration points, and other critical threats to network and user security. Enterprise network managers can use passive monitoring technologies to monitor non-encrypted DNS and utilize standard blocking/filtering/redirection strategies to deal with potentially malicious DNS traffic. DNS traffic over standard ports can also be redirected within the enterprise network to resolvers controlled by the enterprise to ensure policy enforcement and split-DNS integrity.

In an encrypted DNS traffic scenario, enterprise DNS resolvers may still be used to enforce policy and ensure split-DNS integrity, but these passive techniques become more difficult and more expensive to implement. DNS resolution may still be redirected in a DoT scenario given the use of a dedicated port. However, DoH traffic may be more difficult to intercept since it may be commingled with other HTTPS traffic. Enterprises currently deal with HTTPS traffic control/inspection via proxies and other interception technologies, some of those devices may also be able to intercept and block DoH traffic.[40]

Some enterprises perform Man-In-the-Middle (MITM) interception of all HTTPS traffic. This means they decrypt and inspect all HTTPS traffic entering or leaving their network. These enterprises will have an easier time inspecting DoH traffic. Intercepting and decrypting HTTPS traffic may allow for better detection of malcious behavior, but it can also have a significant impact on the security of users and the enterprises that engage in it.[41,42]

Enterprises that do not wish to MITM all HTTPS traffic will not be able to inspect DoH traffic. They may be able to block traffic by either blocking all HTTPS traffic sent to known DoH resolvers, by forbidding the installation of all software that supports DoH, or by managing the configuration of installed software with policies.[43,44,45]

The SSAC notes that RFC 3205 / BCP 56 includes the following recommendation on the use of HTTP as a substrate for other protocols.[46]

> *New protocols - including but not limited to those using HTTP - should not attempt to circumvent users' firewall policies, particularly by masquerading as existing protocols. "Substantially new services" should not reuse existing ports.*

This recommendation in an IETF Best Current Practice publication appears to run counter to the design of DoH. However, privacy necessarily demands some level of concealing both the payload of a transaction from third party inspection, the identity of parties to a transaction, and even the observation that such a transaction took place. This necessity exposes an obvious tension between overt identification of application transactions to the network and a desire by applications not to expose their transactions to the network.

[40] See Protecting Organizations in a World of DoH and DoT, https://live.paloaltonetworks.com/t5/Blogs/Protecting-Organizations-in-a-World-of-DoH-and-DoT/ba-p/313171

[41] See See Durumeric, Z., Ma, Z., Springall, D., Barnes, R., Sullivan, N., Bursztein, E., Bailey, M., Halderman, J., Paxson, V., "The Security Impact of HTTPS Interception", NDSS '17, 26 February–1 March, 2017, San Diego, CA, USA, http://dx.doi.org/10.14722/ndss.2017.23456

[42] See Factsheet TLS interception, https://english.ncsc.nl/publications/factsheets/2019/juni/01/factsheet-tls-interception

[43] See Configure Microsoft Edge policy settings on Windows, https://docs.microsoft.com/en-us/DeployEdge/configure-microsoft-edge

[44] See Enforcing policies on Firefox for Enterprise, https://support.mozilla.org/en-US/kb/enforcing-policies-firefox-enterprise

[45] See Understand Chrome policy management, https://support.google.com/chrome/a/answer/9037717?hl=en

[46] See RFC 3205, Section 9.2

Many enterprises deploy what is commonly referred to as 'split-DNS' or 'private DNS'.[47] With split-DNS, a resolver may return a different answer to a query depending on the source of that query. A user located on the network of an enterprise using split-DNS may have certain names that only resolve when using the enterprise's resolver from within their network. Using a different resolver within the enterprise network may result in queries leaking outside of the enterprise.

One example is accessing an intranet web site that is only meant to be viewed within the enterprise. If that user's browser is using a resolver located outside of the enterprise they would not be able to resolve names internal to the enterprise. Further, attempts to resolve internal addresses by a browser configured to use an external resolver will likely lead to leakage of sensitive network configuration information to the external resolver. This introduces a new attack vector where an external resolver could proffer answers to internal network requests that, in turn, expose sensitive enterprise data. For more detailed information on the security implications of leaked queries and name collisions, please refer to the SSAC's previous reports on the subject.[48,49]

The introduction of new DNS transports, and DoH in particular, threatens to upend this model of network control and management. Instead of deploying a relatively lightweight filter at the DNS level, in order to maintain current capabilities companies may need to intercept and inspect all HTTPS traffic. This may add significant overhead and cost to a company’s IT infrastructure without a corresponding increase in capability. They will have to spend that money and time just to keep the capabilities they have today. They may also determine that it is necessary to enact additional control measures such as decrypting all HTTPS traffic, preventing the configuration of DoH in applications, or preventing the installation of certain applications.

Another option for enterprises is the so called “canary domain”. Mozilla Firefox currently allows network providers to signal that DoH should not be used on their network.[50] Providers can provision the domain **use-application-dns.net** in a specific manner to signal that Mozilla Firefox should not use DoH. This use of a canary domain is currently not standardized and only in use by Mozilla Firefox.

[47] See RFC 8499, section 6

[48] See SAC064

[49] See SAC066

[50] See Canary domain use-application-dns.net, https://support.mozilla.org/en-US/kb/canary-domain-use-application-dnsnet

4.3 Dissidents, Protesters, and Others

The Internet is a valuable means of communication and organization. It enables dissidents and protesters to spread alternative views, critique politics, and shed light on corruption and human rights abuses, activities that are particularly important when state controlled media cannot discuss such content.[51] Dissemination of information over the Internet can also allow organizers of an event to share more relevant information with attendees.[52] The Internet is often the most efficient means to mobilize people and deliver alternative viewpoints.

In addition to government surveillance, some ISP customers believe that their ISPs monetize DNS query data at the expense of their customers, and so prefer to use resolvers run by other entities that they believe will not do that. Other ISP customers believe that their ISPs have effective privacy policies due to local law and practice, and prefer to use their ISP's resolvers. Some users are concerned that some third party resolver providers may be monetizing their DNS queries and prefer a more direct contractual relationship with their DNS resolution provider.

By encrypting DNS queries and resolution, DoH and DoT can help shield users from being tracked by their Internet Service Providers (ISPs). ISPs are heavily influenced or directly controlled by the state in many countries.[53] In some jurisdictions, ISPs are legally required to store users' DNS data for specified amounts of time and to permit intelligence and security agencies to access the data.[54] Thus, ISP tracking can easily facilitate government surveillance, which can be legitimately dangerous for political dissidents. Using a remote DNS resolver shifts the jurisdiction of the query resolution. This may expose users to a different regime's policies, and if there are relevant treaties in-place, merely make it a bit more difficult for local intelligence and security agencies to access the same data.

There are existing privacy tools that match or exceed the technical capabilities of DoH and DoT, such as VPNs and Tor.[55] However, political dissidents often face legal and social barriers when attempting to use them. For example, some countries have placed legal restrictions on the use of VPNs and threaten to fine VPN companies that allow unauthorized access within their

[51] See Ruijgrok, K., "From the web to the streets: internet and protests under authoritarian regimes", Democratization, 24:3, 498-520, 2017, https://doi.org/10.1080/13510347.2016.1223630

[52] *ibid.*

[53] See Christensen, B.,"Cyber state capacity: A model of authoritarian durability, ICTs, and emerging media", Government Information Quarterly, Volume 36, Issue 3, July 2019, 460-468, https://doi.org/10.1016/j.giq.2019.04.004

[54] See Freedom House, Freedom on the Net: The Rise of Digital Authoritarianism, https://freedomhouse.org/report/freedom-net/freedom-net-2018

[55] See The Tor Project, https://www.torproject.org/

jurisdictions.[56] They may also use subpoenas to access VPN companies' records. Social barriers include language differences and lack of technical knowledge.[57] Finally, even though VPNs and Tor can frustrate tracking of dissidents' activities, the use of such tools can, in itself, raise suspicions.[58]

Of course, DNS over encrypted transport is not a panacea. Governments that wish to discourage citizens from communicating and organizing online can often do so by other means, some of which are more intrusive than analyzing DNS data. For example, using DoH might be a trigger to deploy IP tracing and more detailed packet inspection. Some regimes may declare encryption to be illegal or require backdoors in encryption mechanisms.[59] Furthermore, even laws that are difficult or impractical to enforce may cause people to self-censor their online speech.[60] Some may overestimate the regime's capabilities, while others might err on the side of caution if penalties are harsh. Even with DoH or DoT, the ability of citizens to express political dissent without reprisal is greatly influenced by their governments.

The benefits and shortcomings of DoH or DoT largely also apply to visitors from permissive countries who are visiting more restrictive ones. Visitors are unlikely to organize protests online, but their browsing habits and communications may be objectionable to their host governments. Like local dissidents, they have an interest in evading government surveillance. However, governments may find alternative ways to make visitors adhere to their policies. Likewise, advertisers and others interested in monetizing the behavior of Internet users will develop new ways to track user behavior.

4.4 Internet Service Providers

Many national regimes have enacted regulatory measures that oblige ISPs to prevent their customers from accessing certain sites and/or services. Given the common use of shared hosting platforms and content distribution networks, it is uncommon for individual IP addresses to be uniquely associated with individual web sites. Consequently, IP address filtering approaches have a risk of unintended consequences, where access to the restricted site may not necessarily be blocked, and access to many other sites may also be affected. Using DNS-based filtering has

[56] See Freedom House, Freedom on the Net: The Rise of Digital Authoritarianism, https://freedomhouse.org/report/freedom-net/freedom-net-2018

[57] See Nekrasov, M., Parks, L., Belding, E., "Limits to Internet Freedoms: Being Heard in an Increasingly Authoritarian World", LIMITS '17: Proceedings of the 2017 Workshop on Computing Within Limits, 119-128, https://dl.acm.org/citation.cfm?id=3080564

[58] *ibid.*

[59] See Bertola, V., "DNS-over-HTTPS Public Policy Briefing", Section 5.4, Version 1, November 2018, https://www.open-xchange.com/fileadmin/user_upload/Blog/DoH_Public_Policy_Briefing.pdf

[60] See Hellmeier, S., "The Dictators Digital Toolkit: Explaining Variation in Internet Filtering in Authoritarian Regimes", Politics & Policy 44, Issue 6, 1158-1191, https://doi.org/10.1111/polp.12189

proven to be more popular with ISPs, with an implementation based on adding a number of filter rules into the ISP's DNS recursive resolver infrastructure. [61,62,63]

If applications are configured to use a non-local resolver it effectively bypasses the ISP's ability to use the DNS in this manner. The ISP's obligations under the regulatory regime would not be altered, but if an increasing proportion of an ISP's users were to bypass the ISP's DNS resolution infrastructure (such as through applications that use DoH) then the ISP's ability to meet their regulatory requirements may be compromised.

One option for ISPs is to attempt to block all encrypted DNS traffic. In the case of DoT they may be able to argue that blocking all traffic on TCP port 853 is enough to fulfill their regulatory requirements, even if it is possible to run DoT over ports other than 853. However, it will likely not be possible to differentiate DoH traffic from other HTTPS traffic at the scale of an ISP. Thus ISPs will need to block all HTTPS traffic destined for hosts identified as running DoH recursive resolvers. This will undoubtedly miss some DoH traffic destined for servers not identified by the ISP, while also unintentionally blocking some non-DoH HTTPS traffic. Any attempt an ISP will make to block all DoH traffic will involve considerable guesswork.

Many ISPs offer additional services based on DNS filtering for protecting their users, and also their own networks, from malware infiltration and other malicious actors.[64,65] These services function in much the same way as the mechanisms employed in enterprise network filtering discussed in section 4.2. The difference is that they can be turned on or off for specific users of the ISP, often for a price. If users' applications send their DNS queries to a non-ISP DNS resolver these services may no longer work.

If users' applications direct their DNS queries to a non-ISP DNS resolver, this can result in offloading DNS query load from the ISP's resolvers. Some ISPs already do this by pointing their users at existing public DNS resolvers. However, users experiencing problems with DNS will often call their ISP's help desk first to seek assistance. If the user is sending their queries to a DNS server outside of the ISP's control there will be little the ISP can do to help the user, and

[61] See Clayton, R., "Failures in a Hybrid Content Blocking System", International Workshop on Privacy Enhancing Technologies, Springer, Berlin, Heidelberg, 78-92, http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.63.8009&rep=rep1&type=pdf

[62] See Nabi, Z., "The Anatomy of Web Censorship in Pakistan", Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet, 2013, https://www.usenix.org/system/files/conference/foci13/foci13-nabi.pdf

[63] See Levis, P., "The collateral damage of internet censorship by dns injection." ACM SIGCOMM, CCR 42.3, http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.386.7281&rep=rep1&type=pdf

[64] See BT Parental Controls: Keep your children safe online, https://home.bt.com/tech-gadgets/internet/broadband/stay-safe-with-bt-parental-controls-11363887238413

[65] See Internetmatters.org Parental Controls, https://www.internetmatters.org/parental-controls/

this will place an increased technical support burden on the ISP. Simply handling the user's call and helping them understand that the problem lies outside of the ISP's control could consume considerable time of the ISP's call center staff.

5 Application Resolver Choice

Many of the protocols and practices necessary to complete the entirety of a DoH or DoT deployment model are still being designed and implemented. The most important of these is how applications or operating systems discover, or automatically configure, their DoH or DoT resolver. In many networks the Dynamic Host Configuration Protocol (DHCP) is used to configure operating systems with DNS resolver information such as which resolver to use and the existence of local namespaces. However, with DoH and DoT, there is no as yet agreed upon analogue for this functionality provided by DHCP. Users can configure their operating system to use specific DoH or DoT servers, as they usually could with traditional DNS as well, but most users don't and instead rely on automatic configuration and default settings.

Another missing technology is how to signal the evolving relationship between how applications use DNS, how the operating system uses DNS, and how a local network administrator desires users on their network to use DNS. The development of DoH and DoT has led to applications resolving names in ways that may be different than how the operating system resolves names. For example, an application resolving names with DoH may cache names or validate names with DNSSEC while the operating system the application is running on may do neither. At present there is no mechanism for different implementations of DNS resolvers running on the same computer to coordinate their behavior, or even be aware that differences in behavior exist. Some applications had been doing their own DNS resolution prior to the standardization of DoH and DoT, but the expected widespread deployment of applications doing their own DNS resolution will exacerbate this issue.

One option an application designer can make is for their application to choose the same resolver as the operating system. An application using DoH or DoT may be able to inspect the DNS configuration of the operating system, and then try to use the same resolver. However, that resolver may not support DoH or DoT, thereby requiring the application to use the operating system for DNS resolution, or choose a different resolver than what the operating system has configured. Another option is for application designers to choose a specific DoH or DoT resolver without giving consideration to the resolver configured in the operating system.

The introduction of DoH has led to more situations where an application is using a different resolver than the operating system it runs on. This has important implications for users, many of which are discussed in section 4.

6 DNS Resolution in Applications

Traditionally, applications requiring DNS resolution do so through a library that is part of the operating system. This library is configured together with the operating system, and initial configuration is often given by the Internet access provider or a local network administrator. The configuration from the access provider can be static or dynamic, often using DHCP for dynamic configuration, but in both cases the end user can typically, although perhaps not easily, override these settings. The library functions as a stub resolver, and requires a recursive resolver, often on a different computer, to collect DNS responses.

As the SSAC recently wrote in relation to the Internet of Things (IoT),[66] many constrained DNS-capable devices ship without memory capable of storing configuration changes. Their configuration is set at the factory when they are manufactured, and then can never be changed. This hard setting of configuration data ultimately removes control from the user and places it into the hands of device manufacturers, which may result in the device only performing DoH resolution to a pre-configured resolver, or alternatively, result in the device being incapable of being configured to do so.

What is discussed in this report is not only the underlying transport protocol (i.e., HTTPS, TLS), but also how DNS resolution works in practice. DoH implemented in web browsers and other applications moves the initial query from the operating system to the application, which then determines the rest of the resolution path. It has always been possible for any application developer to embed DNS resolution in their application using traditional unencrypted DNS resolution. What is new with DoH is that it is intended to be implemented in applications, thereby bypassing the operating system's DNS resolution path. DoH will likely be the first time that DNS resolution in applications reaches a significant user base by default. DoT can also be implemented in applications, but is more intended to be a replacement for operating system stub resolvers.

Currently, unencrypted DNS over UDP and TCP utilizes a specific port number. This implies that simple Internet traffic flow inspection can detect traditional DNS traffic. When DNS functionality is integrated into applications, these applications no longer use operating system

[66] See SAC105, Section 5.1

provided DNS functions, thereby increasing the difficulty of inspecting DNS traffic flows. It has always been possible for applications to hide traffic in this way, as any protocol can run on any port number, but the default has been to have DNS unencrypted, over the IANA registered UDP and TCP port 53,[67] which makes it relatively easy to detect and inspect it.

In addition to the increased difficulty of monitoring DNS traffic flows, applications performing DNS functions themselves may cause other disruptions which may or may not be visible to users of those applications. For example, DNS implementations in applications might make use of a different set of DNS servers from the normal, publicly visible DNS servers. They may not even be DNS servers as conventionally understood today. An application with its own DNS implementation may or may not reveal to its users that it is using an alternative set of DNS servers. Or applications implementing a DNS function may choose to implement different DNS functionality than is provided by the operating system (e.g., DNSSEC).

One industry concern with respect to applications providing DNS functionality is that they will undermine the usefulness of DNS as a generic, protocol-neutral naming system for the Internet. As originally deployed, DNS was application agnostic because it was configured per host, not per application or per user protocol. There are many record types (both RRtypes and EDNS options) that are defined for the use of specific protocols, but the general resolution mechanism and the basic implementation of the DNS is not specific to either who is resolving a name, or why. Many Internet applications depend on the DNS, not just the web, and application agnostic DNS has been an important principle in the Internet's development. Content Distribution Networks (CDNs) and other services have bent this principle, but optimizing DNS for HTTPS seems likely to bend it further, or possibly even render it irrelevant. Tailoring DNS specifically for HTTPS may interfere with the permissionless innovation that has allowed the Internet to be so successful.

Another industry concern is that DoH has the potential to change the relationships between applications and common infrastructure. The traditional mode of interaction between applications and the namespace is via a set of library calls provided by the host platform that interact with resolvers provisioned via the ISP or enterprise. Any policies or practices associated with the resolution of names can be implemented via the ISP or enterprise. At the same time, the end user provides an implicit data stream to the ISP or enterprise via the DNS query stream. An application using DoH has the ability to select its own DNS name resolvers, and communicate with them in a manner which is, at a first level, indistinguishable from any other HTTPS data flow. The application has the ability to select a resolver according to its own preferences and may choose to avoid using not only the common local host's name resolution services, but also

[67] See IANA Service Name and Transport Protocol Port Number Registry, https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml

the ISP's or enterprise's name resolution services. This opens up the possibility for application-specific namespaces that are tailored to the requirements of a particular application, and not generally visible to any other component of the Internet. This shift in the fabric of the namespace could undermine the utility and value of a common protocol and provider agnostic namespace for the Internet.

In the past, web browsers cached identifiers globally, allowing names to be shared between different sites accessing the same content. Content that was downloaded for one origin (i.e., per-browser tab) was cached globally in the browser and available in the cache for the browser to use when rendering other websites. Recently due to security concerns, browsers have begun changing to caching identifiers per-origin (i.e., per-website, per-tab).[68,69,70,71] With this change, browsers may download the same web resource twice if it is used by two separate sites. This also means that, in practice, each website will have its own cached versions of content. Content accessed at one time may not be the same when it is accessed a second time for a different website. In the future with DNS resolution in the browser, it may no longer be necessary for each website to even use the same namespace.

7 Conclusions

Evaluations of DoH or DoT rely on the perspective of the evaluator. Thus any comprehensive understanding of the technologies requires an appreciation for the implications of the technology on different user groups and their perspectives. In addition, how DoH or DoT are implemented in applications and operating systems, and how these are deployed in networks, will affect each user group differently. DoH and DoT are just protocols for transporting DNS queries and responses. How they are implemented, how they are deployed, what default settings are configured, and who uses them, are the questions that this report focuses on.

The discussions in section 4 did not include discussions about the differences between the choice of DoH or DoT and how they might affect different parties who use or manage the DNS. This approach was deliberate, as both DoH and DoT have very similar properties, with the commonly debated topics being related to the implementation or policies around the protocols.

[68] See Chrome Platform Status, Partition the HTTP Cache, https://www.chromestatus.com/feature/5730772021411840

[69] See Mozilla Bugzilla, Partition the HTTP cache per the top-level document's origin, https://bugzilla.mozilla.org/show_bug.cgi?id=1536058

[70] See WebKit Bugzilla, Optionally partition cache to prevent using cache for tracking, https://bugs.webkit.org/show_bug.cgi?id=110269

[71] See Double-keyed HTTP cache, https://github.com/whatwg/fetch/issues/904

DoH and DoT are very similar, but not identical. As an example, a network operator may wish to block encrypted DNS traffic to some endpoints or to capture all DNS traffic in order to apply business logic or filtering. With DoH, this is very difficult in the default state for cases where the DoH endpoint is shared infrastructure. In the case of DoT this issue would not be as prominent when the default settings are used. However, the specification for DoT allows for the port to be changed, possibly to port 443 if that is mutually agreed upon by the server and client, which could be owned and operated by the same entity.

Another discussion area around encrypted transport protocols is how the queries are directed. The specifications of DoH and DoT do not define which types of queries are sent to local resolvers or to other providers. Instead those logistical details are controlled by the system or application that is issuing the queries, and in some but not all cases, these configuration values are exposed to the end users.

While the introduction of these encrypted transport methods for DNS can provide point to point secrecy for DNS queries, their policies and implementations can dramatically change where user query information is handled and who has the ability to act on the data in transit, with the potential for both beneficial and harmful results.

As DoH and DoT continue to be implemented, deployed, configured, and used, the perspectives listed in Section 4 of this report will change. Some users will find that their initial concerns regarding the technology were unfounded, while others will experience as-yet undocumented issues with specific deployments. Some network administrators will discover novel methods to circumvent or block DoH to meet their requirements, and users insistent on using DoH will circumvent these circumventions. In short, the Internet will continue to evolve.

Regardless of perspective, the deployment of DoT, and particularly DoH will be disruptive. The disruption lies mainly in the implementation and deployment of the technology, with application-specific DNS resolution via a commingled encryption stream providing a host of challenges and changed assumptions about how networks and endpoints work, who has access to DNS resolution data, and how to protect and manage networks in this new model.

8 Acknowledgments, Statements of Interest, and Dissents and Withdrawals

In the interest of transparency, these sections provide the reader with information about four aspects of the SSAC process. The Acknowledgments section lists the SSAC members, outside experts, and ICANN staff who co-authored or contributed directly to this particular document

(Contributors) or who provided reviews (Reviewers). The Statements of Interest section points to the biographies of all SSAC members, which disclose any interests that might represent a conflict—real, apparent, or potential—with a member's participation in the preparation of this Report. The Dissents section provides a place for individual members to describe any disagreement that they may have with the content of this document or the process for preparing it. The Withdrawals section identifies individual members who have recused themselves from discussion of the topic with which this Report is concerned. Except for members listed in the Dissents and Withdrawals section, this document has the consensus approval of all of the members of SSAC.

8.1 Acknowledgments

The committee wishes to thank the following SSAC members for their time, contributions, and review in producing this report.

Contributors
Barry Leiba (work party co-chair)
Suzanne Woolf (work party co-chair)
Joe Abley
Tim April
Paul Ebersman
Ondrej Filip
Geoff Huston
Warren Kumari
Jacques Latour
John Levine
Chris Roosenraad
Tara Whalen

Reviewers
Greg Aaron
Jaap Akkerhuis
Lyman Chapin
KC Claffy
Patrik Fältström
James Galvin
Andrei Kolesnikov
Carlos Martinez
Danny McPherson
Ram Mohan

Observers
Julie Hammer
Merike Kaeo
Rod Rasmussen

ICANN staff
Patrick Jones
Andrew McConachie (editor)
Danielle Rutherford
Kathy Schnitt
Steve Sheng

8.2 Statements of Interest

SSAC member biographical information and Statements of Interest are available at: https://www.icann.org/resources/pages/ssac-biographies-2019-11-20-en

8.3 Dissents and Withdrawals

There were no dissents or withdrawals.

Appendix A: Glossary

Authoritative Server

A server that knows the content of a DNS zone from local knowledge, and thus can answer queries about that zone without needing to query other servers.

RFC 2182

Bring Your Own Device (BYOD)

A term that refers to being allowed to use one's personally owned device, rather than being required to use an officially provided device. Typically it refers to employees using their personal devices where they work.

https://en.wikipedia.org/wiki/Bring_your_own_device

Classic DNS

DNS over UDP or TCP as defined in RFC1035 and its successors. It applies to DNS communication between stub resolvers and recursive resolvers, and between recursive resolvers and authoritative servers. It is not encrypted. Also called *Traditional DNS*.

draft-ietf-dnsop-terminology-ter-01, https://datatracker.ietf.org/doc/draft-ietf-dnsop-terminology-ter/

DNS Blocking

A strategy for making it difficult for users to resolve specific domains by blocking or synthesizing responses to queries. Also called DNS filtering. Also see Response Policy Zone.

https://en.wikipedia.org/wiki/DNS_blocking

DNS Firewall

A firewall that is capable of blocking network traffic based on information learned by inspecting DNS traffic.

DNS over HTTPS (DoH)

A protocol for sending DNS queries and getting responses over Hyper Text Transfer Protocol (HTTP) using Secure Hyper Text Transfer Protocol (HTTPS) (and therefore Transport Layer Security (TLS) for integrity and confidentiality). Each DNS query-response pair is mapped into an HTTP exchange.

RFC 8484

DNS over TLS (DoT)

A protocol for sending DNS queries and getting responses over Transport Layer Security (TLS) on TCP port 853.

RFC 7858

Domain Name System Security Extensions (DNSSEC)

Extensions to the DNS protocol that add data origin authentication and data integrity verification.
RFC 4033

Dynamic Host Configuration Protocol (DHCP)

A protocol for providing configuration parameters to hosts. DHCP is often used to automatically configure the DNS recursive resolver of hosts.
RFC 2131

Firewall

A network security system that monitors and controls incoming and outgoing network traffic based on predetermined security rules.
https://en.wikipedia.org/wiki/Firewall_(computing)

Encrypted DNS

DNS over encrypted transport, including both DNS-over-TLS and DNS-over-HTTPS.

Internet Service Provider (ISP)

An organization that provides services for accessing, using, or participating in the Internet.
https://en.wikipedia.org/wiki/Internet_service_provider

Man in the Middle (MITM)

A technique whereby traffic is decrypted in transit and inspected or altered. Typically this happens transparently, without either end-point being aware of the event. MITM can refer to either the actor performing the technique or the act itself.
https://en.wikipedia.org/wiki/Man-in-the-middle_attack

Query Name (QNAME) Minimization

A technique to improve DNS privacy where a DNS resolver does not send the original full queried name to authoritative name servers.
RFC 7816

Quick UDP Internet Connections (QUIC)

A secure general-purpose transport protocol that uses the User Datagram Protocol (UDP). QUIC authenticates all of its headers and encrypts most of the data it exchanges.
draft-ietf-quic-transport-27, https://datatracker.ietf.org/doc/draft-ietf-quic-transport/

Recursive Resolver

A DNS resolver that performs queries on behalf of a stub resolver and caches the answers it receives.
RFC 8499

Response Policy Zone (RPZ)

A mechanism to introduce a customized policy in DNS servers, so that recursive resolvers return possibly modified results. By modifying a result, access to the corresponding host can be blocked. Also see DNS Blocking.
https://en.wikipedia.org/wiki/Response_policy_zone

Secure HyperText Transfer Protocol (HTTPS)

A protocol for sending Hyper Text Transfer Protocol (HTTP) over Transport Layer Security (TLS).
RFC 2818

Stub Resolver

A resolver that cannot perform all resolution itself and typically depends on a recursive resolver to undertake the actual resolution function.
RFC 8499

Tor

A protocol with associated software and network of Tor nodes that provides encrypted transport and traffic anonymization.
https://www.torproject.org/

Traditional DNS

See *Classic DNS*.

Transmission Control Protocol (TCP)

A connection oriented protocol that is one of the main protocols of the Internet protocol suite.
https://en.wikipedia.org/wiki/Transmission_Control_Protocol

Transport Layer Security (TLS)

A protocol that allows client/server applications to communicate over the Internet in a way that is designed to prevent eavesdropping, tampering, and message forgery.
RFC 8446

User Datagram Protocol (UDP)

A connectionless protocol that is one of the main protocols of the Internet protocol suite.
https://en.wikipedia.org/wiki/User_Datagram_Protocol

Virtual Private Network (VPN)

A generic term that covers the use of public or private networks to create groups of users that are logically separated from other network users and that may communicate among themselves as if they were on a private network.
RFC 4026

Appendix B: Additional Resources

This appendix contains additional resources that the SSAC found informative and useful when preparing this report. Many of the technologies discussed in this report are changing rapidly. The snapshot of resources listed in this section represents the state of the art at publication time of this report, and may have since been superseded by new developments.

Böttger, T., Cuadrado, F., Antichi, G., Fernandes, E.L.,, Tyson, G., Castro, I., Uhlig, S., "An Empirical Study of the Cost of DNS-over-HTTPS", In Internet Measurement Conference (IMC ’19), October 21–23, 2019, Amsterdam, Netherlands. https://arxiv.org/abs/1909.06192

Chromium Blog, "Addressing some misconceptions about our plans for improving the security of DNS", https://blog.chromium.org/2019/10/addressing-some-misconceptions-about.html

Council of European National Top-Level Domain Registries, "CENTR Issue Paper on DNS over HTTPs",
https://centr.org/library/library/policy-document/centr-issue-paper-on-dns-over-https.html

Deccio, C., Davis, J., "DNS privacy in practice and preparation", CoNEXT '19: Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, December 2019 Pages 138–143, https://doi.org/10.1145/3359989.3365435

ICANN Office of the CTO, "OCTO-003v2: Local and Internet Policy Implications of Encrypted DNS", Paul Hoffman, https://www.icann.org/en/system/files/files/octo-003-en.pdf

The Internet Society, "Consolidation in the Internet Economy", 2019,
https://future.internetsociety.org/2019/

Lu, C., Liu, B., Li, Z., Hao, S., Duan, H., Zhang, M., Leng, C., Liu, Y., Zhang, Z., Wu, J., "An End-to-End, Large-Scale Measurement of DNS-over-Encryption: How Far Have We Come?", In Internet Measurement Conference (IMC ’19), October 21–23, 2019, Amsterdam, Netherlands. https://doi.org/10.1145/3355369.3355580

The Messaging, Malware and Mobile Anti-Abuse Working Group, "Tutorial on Third Party Recursive Resolvers and Encrypting DNS Stub Resolver-to-Recursive Resolver Traffic", https://www.m3aawg.org/dns-crypto-tutorial

The Messaging, Malware and Mobile Anti-Abuse Working Group, "Companion Document: Recipes for Encrypting DNS Stub Resolver-to-Recursive Resolver Traffic", https://www.m3aawg.org/dns-crypto-recipes

Microsoft Windows, Tech Net, "Windows will improve user privacy with DNS over HTTPS", https://techcommunity.microsoft.com/t5/Networking-Blog/Windows-will-improve-user-privacy-with-DNS-over-HTTPS/ba-p/1014229

Rescorla, E., "Status of DoH/TRR in Firefox", In Domain Name System Operations Analysis and Research Center 32, February 8-10, 2020, San Francisco, USA. https://indico.dns-oarc.net/event/33/contributions/750/

National Institute of Standards, "NIST Special Publication 800-81-2, Secure Domain Name System (DNS) Deployment Guide", http://dx.doi.org/10.6028/NIST.SP.800-81-2

RIPE Labs, "DNS Censorship (DNS Lies) As Seen By RIPE Atlas", Stéphane Bortzmeyer, https://labs.ripe.net/Members/stephane_bortzmeyer/dns-censorship-dns-lies-seen-by-atlas-probes