

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 25, 2019

P. Hoffman
ICANN
March 24, 2019

Associating a DoH Server with a Resolver
draft-ietf-doh-resolver-associated-doh-03

Abstract

Browsers and web applications may want to know if there are one or more DoH servers associated with the DNS recursive resolver that the operating system is already using. This would allow them to get DNS responses from a resolver that the user (or, more likely, the user's network administrator) has already chosen. This document describes two protocols for a resolver to tell a client what its associated DoH servers are. It also describes a protocol for a client to find out the address of the resolver it is using, if it cannot find that address by an operating system API or some other means.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 25, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	4
2. DoH Servers from HTTPS	4
3. DoH Servers from DNS	5
4. Resolver Addresses from DNS	6
5. Lists of DoH Servers	7
6. IANA Considerations	7
7. Privacy Considerations	7
8. Security Considerations	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Appendix A. Earlier Design Choices	10
Acknowledgments	10
Author's Address	10

1. Introduction

DoH [RFC8484] requires that one or more DoH servers be configured for the DoH client. That document does not say how the DoH servers are found, nor how to select from a list of possible DoH servers, nor what the user interface (UI) for the configuration should be.

There is a use case for browsers and web applications to want the DNS recursive resolver(s) configured in the operating system to use DoH for DNS resolution instead of normal DNS, but to do so to at a DoH server specified by the configured resolver. For example, a recursive resolver configured by the operating system may know how to give correct answers to DNS queries that contain names that are only resolvable in the local context, or resolve differently in the local context. Similarly, the recursive resolver configured in the operating system may implement security policies such as malware prevention that are not implemented in the same way in DoH servers not affiliated with the user's organization. Users typically configure their DNS recursive resolvers with through automatic configuration from a protocol such as DHCP; much less often, they use manual configuration (such as manually editing a `/etc/resolv.conf` file).

The expected use cases for DoH are browsers and web applications that would otherwise get their DNS service from the resolver configured by

the operating system. The user of the client might have a preference for using a DoH server for the benefits that DoH brings, and they might need to use a DoH server that is associated with the resolver that the computer is currently using for the reasons listed above. In a common scenario, user may be required to use only resolvers that are approved by their organization's network operators.

The URI templates of the DoH servers associated with a resolver might be hosted on the resolver itself, or a resolver hosted by the same operator, or even hosted somewhere else. The latter could be used by resolver operators who don't want to host DoH servers but trust another operator to do so.

To address these use cases, this document defines protocols to get the list of URI templates [RFC6570] or addresses for the DoH servers associated with at least one of the resolvers being used by the operating system on the system on which the application is being run.

- o "DoH servers from HTTPS", described in Section 2, is a well-known URI [I-D.nottingham-rfc5785bis] that can be resolved to return the URI templates in an HTTP response.
- o "DoH servers from DNS", described in Section 3, is a new special use domain name (SUDN) [RFC6761] that can be queried to return the URI templates as a TXT RRset.
- o "resolver addresses from DNS", described in Section 4, is a new SUDN that that can be queried to return the addresses as A and AAAA RRsets.

For these protocols to be useful in a browser, the browser needs to have an entry in its configuration interface where the allowed DoH servers are listed that indicates that a DoH server from the configured Do53 or DoT resolver is allowed. That wording might say something like "DoH server associated with my current resolver" (or "servidor DoH asociado con mi resolucioⁿ actual" or "serveur DoH associe a mon resolveur actuel"). Alternatively, these protocols might be the default for a browser, and naming specific other DoH servers might be done with a UI.

The protocols described here are meant for a browser to be able to start using DoH based on its current interactions with the resolver from the operating system on which the browser is running: they are not expected to work without being able to do DNS resolution. Even "DoH servers from HTTPS", which ostensibly only needs an IP address, is likely to need DNS resolution for things like OCSP servers during the setup of TLS.

1.1. Terminology

In this document, "client" means either a browser or web application. When one or the other is named explicitly,

In this document, "browser" means any application that can open ports in the operating system. This odd usage of the term "browser" is adopted from [RFC8484].

In this document, "DoT" is used to indicate DNS over TLS as defined in [RFC7858].

In this document, "Do53" is used to indicate DNS over UDP or TCP as defined in [RFC1035].

"DoH client" and "DoH server" are defined in [RFC8484].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. DoH Servers from HTTPS

To find the URI templates for DoH servers associated with a resolver whose address is already known, a browser or web application can use the well-known URI described in this section. To find the IP address of the resolver used by the operating system on which a browser is running, a browser can use either an operating system function (if such a function is available to it) or the process described in Section 4. (A web application can also use the process described in Section 4, and thus use this well-known URI.)

To find the DoH servers associated with a resolver, the client MUST use the following query:

`https://IPADDRESSGOESHERE/.well-known/doh-servers-associated/`

The resolver replies with its associated DoH servers as URI templates [RFC6570]. The HTTP response header MUST contain an appropriate HTTP status code as defined in [RFC7231]. Successful responses MUST set the Content-Type set to "application/json".

The returned JSON object [RFC8259] MUST contain a member whose name is "associated-resolvers" and whose value is a JSON array. The array contains zero or more JSON strings, each of which is a single URI template.

If the array of URI templates returned is empty, that indicates that the resolver does not have any DoH servers associated with it.

For example, the returned object might look like:

```
{ "associated-resolvers":  
  [ "https://dnsserver.example.net/dns-query{?dns}",  
    "https://webhost.example.net/a/b/c/dns-query{?dns}" ]  
}
```

If there are no associated DoH servers, the returned object would look like:

```
{ "associated-resolvers": [ ] }
```

If the TLS authentication for the query fails, the browser MUST abort the connection without sending the HTTP request, and it cannot assume anything about whether the resolver has any DoH servers associated with it.

A client using this protocol MUST try to establish a new list of DoH servers associated with a resolver every time the configured resolver in the operating system changes.

The HTTP query follows all the normal rules for HTTP. Thus, the result of sending this query can be an HTTP redirect to a different server. Also, the result of the query might be served from an HTTP cache.

3. DoH Servers from DNS

To find the URI templates for DoH servers associated with a resolver, a browser sends that resolver a query for "resolver-associated-doh.arpa" in class IN with the RRtype of TXT [RFC1035] (that is, the query is resolver-associated-doh.arpa/IN/TXT). This protocol is most likely useful only to browsers that can call operating system functions that in turn query the DNS for text records. Web applications cannot currently call such operating system functions.

As described in Section 6, the zone resolver-associated-doh.arpa is not actually delegated and never will be. The resolver that receives this query acts as if it is delegated, and adds its own TXT records to the answer. The resolver replies with its associated DoH servers as URI templates in the TXT RRset in the Answer section. The resolver can generate this reply with special code to capture queries for "resolver-associated-doh.arpa"; if the resolver can be configured to also be authoritative for some zones, it can use that

configuration to actually be authoritative for "resolver-associated-doh.arpa".

A resolver that understands this protocol MUST send a TXT RRset in the Answer section. Each TXT record contains one URI template. If a resolver that understands this protocol has no associated DoH servers, the TXT RRset contains exactly one record that has an empty string as the RDATA; that is, the RLENGTH in that record is 1, and the RDATA contains just the byte 0x00.

An example of the TXT RRset (in DNS master file format) might be:

```
$ORIGIN resolver-associated-doh.arpa.  
IN TXT "https://dnsserver.example.net/dns-query{?dns}"  
IN TXT "https://webhost.example.net/a/b/c/dns-query{?dns}"
```

If there are no associated DoH servers, an example of the the TXT RRset (in DNS master file format) might be:

```
$ORIGIN resolver-associated-doh.arpa.  
IN TXT ""
```

The client uses the TXT records in the response to the resolver-associated-doh.arpa/IN/TXT query as a list of the URI templates of the DoH servers associated with the resolver. Note that TXT records can contain multiple "character-strings" [RFC1035]; for this protocol, all characters-strings in a TXT record are concatenated to form a single URI template.

A client using this protocol MUST try to establish a new list of DoH servers associated with a resolver every time the configured resolver in the operating system changes.

4. Resolver Addresses from DNS

Browsers which cannot get the IP address(es) of the resolver configured by the operating system using APIs are still able to use an operating system function such as `gethostbyname()` or its equivalents to convert host names into IP addresses through the stub resolver in the operating system on which they are running. Web applications also can convert host names to IP addresses. Either can use a new SUDN to find the address(es) of the resolvers configured by the operating system.

A browser or web application uses its normal interface for getting IP addresses for a hostname, and uses the SUDN "resolver-addresses.arpa" as the hostname.

As described in Section 6, the zone resolver-addresses.arpa is not actually delegated and never will be. The resolver acts as if that name is delegated, and returns its own A or AAAA addresses in the records in the answer. The resolver can generate this reply with special code to capture queries for "resolver-addresses.arpa"; if the resolver can be configured to also be authoritative for some zones, it can use that configuration to actually be authoritative for "resolver-addresses.arpa".

A client using this protocol MUST try to establish a new list of DoH servers associated with a resolver every time the configured resolver in the operating system changes.

5. Lists of DoH Servers

The "DoH servers from HTTPS" "DoH servers from DNS" return lists of DoH servers, and those lists can have more than one element. The DoH client can choose any of the servers in the list; that is, there is no inherent "preference" for any of the servers returned. From a mathematical viewpoint, the lists can better be considered as sets.

6. IANA Considerations

IANA will record the domain name "resolver-associated-doh.arpa" in the "Special-Use Domain Names" registry [SUDN]. IANA MUST NOT delegate resolver-associated-doh.arpa in the .arpa zone.

IANA will record the domain name "resolver-addresses.arpa" in the "Special-Use Domain Names" registry [SUDN]. IANA MUST NOT delegate resolver-addresses.arpa in the .arpa zone.

Before this draft is complete, mail will be sent to wellknown-uri-review@ietf.org in order to be registered in the "Well-Known URIs" registry at IANA. The mail will contain the following:

URI suffix: doh-servers-associated
Change controller: IETF
Specification document(s): draft-ietf-doh-resolver-associated-doh
Status: permanent

7. Privacy Considerations

Allowing a user to use DoH to a server associated with the resolver in use by the operating system on the user's machine, instead of using Do53 to the resolver in use by the operating system on the user's machine, can increase communication privacy because of the TLS protection. However, using a DoH server can also reduce overall

privacy because both TLS and HTTPS allow for user identification in ways that plain Do53 does not.

When a Do53 or DoT server indicates that a particular DoH server is associated with it, the client might assume that the DoH server has the same information privacy policies as the Do53 or DoT server. Therefore, a Do53 or DoT server SHOULD NOT recommend a DoH server unless that DoH server has the same (or better) information privacy policy as the Do53 or DoT server.

A browser that has both a stub resolver stack and a TLS stack that is independent of HTTP could make a DoT connection to the resolver being used by the operating system.

8. Security Considerations

If DNS queries sent from stub resolvers to recursive resolvers are not sent over transports that assure data integrity and server authentication, the "DoH servers from DNS" and "Resolver addresses from DNS" protocols are susceptible to on-path attackers directing a user to a DoH server that is not actually associated with their resolver. Do53 is not a secure transport, and neither is DoT using the opportunistic profile.

The DNS responses used in "DoH servers from DNS" and "Resolver addresses from DNS" cannot be validated with DNSSEC [RFC4033], and thus even a validating stub resolver would treat them the same as any other DNS responses in unsigned zones.

There is currently no way for an application to know whether the operating system's stub resolver is using a transport that assures data integrity such as DoT. Even if an application could determine the use of a transport like DoT, the application would also need to know whether the transport was authenticated or was simply chosen opportunistically.

9. References

9.1. Normative References

- [I-D.nottingham-rfc5785bis]
Nottingham, M., "Well-Known Uniform Resource Identifiers (URIs)", draft-nottingham-rfc5785bis-09 (work in progress), February 2019.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", RFC 6570, DOI 10.17487/RFC6570, March 2012, <<https://www.rfc-editor.org/info/rfc6570>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [SUDN] "Special-Use Domain Names", n.d., <<https://www.iana.org/assignments/special-use-domain-names/>>.

9.2. Informative References

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/info/rfc6761>>.
- [RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <<https://www.rfc-editor.org/info/rfc6891>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.

[RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.

Appendix A. Earlier Design Choices

The primary use case for these protocols is a browser or web application that is getting name resolution through the stub resolver on the computer on which it is running wanting to switch its name resolution to DoH.

An earlier design suggestion was to use a new RRtype with a query to `./IN/NEWRRTYPE`. However, it was pointed out that this would not work going through stub resolvers that validate DNSSEC.

An earlier design suggestion was to use DHCP to tell the operating system the DoH servers that the stub resolver might use. That protocol is orthogonal to the one in this document in that it addresses a different use case. If both the protocol in this document and a DHCP-based protocol are standardized, they could co-exist. However, there is no current mechanism for a stub resolver to tell a browser, or a web application, what DoH server the stub resolver is using, so DoH configuration in the stub resolver would not prevent the browser from trying to find a DoH server on its own.

An earlier design suggestion was to use an EDNS0 [RFC6891] extension. The design chosen in this document meets the use case better because applications cannot communicate EDNS0 extensions to the stub resolver.

Acknowledgments

The use case in this document was inspired by discussions and the DRIU BoF at IETF 102 and later in the DNSOP Working Group. Vladimir Cunat, Philip Homburg, Shumon Huque, Martin Thomson, Eric Rescorla, and Tony Finch offered useful advice to improve versions of the protocol before it came to the DOH Working Group.

Author's Address

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org