

Associating a DoH Server with a Resolver
draft-hoffman-resolver-associated-doh-07

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 June 17, 2019.

Copyright Notice

Copyright (c) 2018 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.	Finding the URI Templates of the DoH Servers Associated with a Resolver	4
2.1.	DoH Servers by TXT	4
2.2.	DoH Servers by Addresses	5
2.3.	Issues Common to "DoH Servers by TXT" and "Resolver Addresses by SUDN"	6
3.	Finding the Resolver Addresses Without Operating System APIs	6
4.	User Interface	6
5.	Design Choices	7
6.	IANA Considerations	7
7.	Privacy Considerations	8
8.	Security Considerations	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	9
	Acknowledgments	9
	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

Hoffman

Expires June 17, 2019

[Page 2]

manual configuration (such as manually editing a `/etc/named.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.

To address these use cases, this document defines two protocols to get the list of URI templates [[RFC6570](#)] 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 The first, called "DoH servers by TXT" and described in [Section 2.1](#), is a new special use domain name (SUDN) [[RFC6761](#)] that can be queried for a TXT RRset. 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 can only query for IP addresses.
- o The second, called "DoH servers by addresses" and described in [Section 2.2](#), is a well-known URI [[I-D.nottingham-rfc5785bis](#)] that can be resolved to return the URI templates. This is useful if a browser can call operating system functions that will return the address of the recursive DNS resolver that the operating system is currently using.

This document also defines a third protocol, called "resolver addresses by SUDN" and described in [Section 3](#), that is a new SUDN that that can be queried for the IP address(es) of a resolver. This protocol is useful for a client that can query for the addresses associated with a domain name (such as using the POSIX `"getaddrinfo()"` function) but cannot use an operating system function to find those addresses. For browsers, it is only needed if the browser cannot use an API to determine the configured resolver IP address(es).

The design choices for this protocol, particularly earlier designs that were deemed unusable, are described in [Section 5](#).

1.1. Terminology

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

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. Finding the URI Templates of the DoH Servers Associated with a Resolver

A client (a browser or web application) uses either the protocol in [Section 2.1](#) or [Section 2.2](#) to get a list of URI templates for the DoH servers associated with a resolver. The following sub-sections describe the protocols and have notes that are common to both protocols.

2.1. DoH Servers by TXT

To find the DoH Servers associated with a resolver, an application 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).

As described in [Section 6](#), the zone resolver-associated-doh.arpa is not actually delegated and never will be. The resolver 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 RDLENGTH in that record is 1, and the RDATA contains just the byte 0x00.

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 character-strings in a TXT record are concatenated to form a single URI template.

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.

2.2. DoH Servers by Addresses

To find the DoH servers associated with a resolver, a browser or web application uses either an operating system function (if such a function is available to it) or the process described in [Section 3](#) to find one or more IP addresses for the resolver. It uses one or more of those IP addresses as part of a well-known URI to find out the URI templates [RFC6570] to use for the DoH server(s) associated with the resolver.

To find the DoH servers associated with a resolver, the client sends a query to

<https://IPADDRESSGOESHERE/.well-known/doh-servers-associated/>

The resolver replies with its associated DoH servers as URI templates [RFC6570].

[[Need to describe the media types; likely JSON]]

[[Need to talk about what a response with an empty list means]]

[[Need to talk about what happens if authentication fails. This is complicated by the fact that the application doesn't know if the OS-to-resolver communication is authenticated.]]

[[Need to talk about HTTP caching]]

A client 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 result of sending this query can be an HTTP redirect to a different server, such as when a resolver operator doesn't want to run its own DoH server.

2.3. Issues Common to "DoH Servers by TXT" and "Resolver Addresses by SUDN"

See [Section 8](#) for warnings about sending the DNS queries over a transport that does not assure data integrity (such as Do53), and over a transport that does assure data integrity (such as DoT) but in circumstances where the browser or web application doesn't know the type of DNS transport being used.

A client MUST re-issue the queries in {#doh_by_txt} and {#resolver_by_sudn} every time the configured resolver in the operating system changes.

[[What if there is a list of DoH servers? Pick one (how?) or jump around?]]

3. Finding the Resolver Addresses Without Operating System APIs

Browsers can often, but not always, get the IP address(es) of the resolver configured by the operating system using APIs. Browsers which cannot 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 protocol to find the address(es) of the resolvers configured by the operating system.

In this protocol, the browser or web application uses its normal interface for getting 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".

4. User Interface

For this protocol 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").

5. Design Choices

The primary use case for this protocol 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. A secondary use case is an operating system that wants to make a similar switch.

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.

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-hoffman-resolver-associated-doh](#)

(or its successor, if it is adopted in a WG)

Status: permanent

7. Privacy Considerations

Allowing a user to use DoH instead of Do53 increases communication privacy because of the TLS protection.

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

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-08](#) (work in progress), October 2018.

[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

- [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>>.

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 early versions of the protocol.

Author's Address

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org