

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 3, 2019

M. Nottingham
July 02, 2018

DOH Digests
draft-nottingham-doh-digests-00

Abstract

The lack of flexible configuration and selection mechanisms for DOH servers is identified as suboptimal for privacy and performance in some applications.

This document makes a straw-man proposal for an improvement.

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 January 3, 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.	DOH's Additional Benefits for Associated Services	2
1.2.	Achieving DOH's Privacy Goals through Diversity	3
2.	Conventions and Definitions	4
3.	DOH Digests	4
3.1.	Using DOH Digests	4
3.2.	The DOH Digest Format	5
3.3.	Hostname Normalisation	5
4.	Security Considerations	5
5.	IANA Considerations	5
6.	References	5
6.1.	Normative References	5
6.2.	Informative References	6
	Author's Address	6

[1.](#) Introduction

One of the core motivations for DOH [[I-D.ietf-doh-dns-over-https](#)] is to improve end-user privacy by obfuscating the stream of DNS requests that the DOH client makes. It does this by mixing DOH requests into a stream of "normal" HTTP requests to a configured Web server; for example, a large Web site or a Content Delivery Network.

However, DOH intentionally avoids defining a mechanism for configuring a particular DOH server for a given application or host. So far, the most common way to do so is to select one from a pre-configured list of services in an application, such as a Web browser.

Typically, the list of available DOH services is vetted by the application's vendor to assure that they will honour the application's requirements for handling of sensitive data (i.e., the client's DNS request stream) and similar concerns.

This document proposes a means of selecting a DOH server that encourages the deployment of DOH servers by sharing some of its additional benefits with servers that are good candidates for serving DOH traffic.

[1.1.](#) DOH's Additional Benefits for Associated Services

When a DOH server is colocated with (or closely coordinated with) other network services - especially HTTP services - those associated services enjoy a few additional benefits beyond those seen by adopting DOH in the first place.

- o Associated services have an additional privacy benefit; there is one less party involved in the interaction, whereas "normal" DNS and DOH to an unassociated HTTP server require a third party to resolve names.
- o Removing a third party also removes a separate point of potential failure, improving control over service quality and availability. See [[fragile](#)] for further discussion.
- o Finally, the DOH server can use DNS to optimise the provision of associated services. For example, DNS results can be optimised based on the client's request stream with a higher degree of certainty.

In the future, a DOH server might use Secondary Certificates [[I-D.ietf-httpbis-http2-secondary-certs](#)] to further optimise performance of associated services, by using the information in the DNS request stream to aggregate all of its traffic into a small number of connections (possibly only one), thereby allowing greater coordination of congestion control and avoiding connection setup costs.

1.2. Achieving DOH's Privacy Goals through Diversity

Overall, a major goal for deployment of DOH is to assure that DNS connectivity is robust and private. Arguably, this is best served by having a diverse set of available DOH servers that are colocated with popular HTTP content, so that it's more difficult to discriminate DOH from "regular" HTTP, and so the it's more difficult to block DOH services, due to the high impact of blocking a popular site.

One way to encourage the development of such a set is to offer the additional benefits above to parties that are good candidates for serving such traffic. When clients can direct their DOH queries to the HTTP server which will eventually serve their traffic, it provides both better privacy properties and better performance and availability to a broader set of servers.

This is a marked improvement over the static configuration mechanism commonly in place now; accruing such privacy, availability, and performance benefits to whatever DOH server the application or user selects means that only parties who have a relationship with that service will realise these benefits.

This document proposes one way to achieve this.

2. Conventions and Definitions

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.

3. DOH Digests

A DOH Digest is a Bloom filter indicating the set of hosts a given DOH server should be used for.

3.1. Using DOH Digests

When an application has a valid DOH digest for a given DOH server, it tests the digest for each DNS request it makes by hostname; if the hostname (after normalisation) is found in the digest, all DNS requests regarding that hostname SHOULD be sent to the corresponding DOH server. If multiple DOH digests match a given hostname, any matching DOH server MAY be used; the client SHOULD select one of the candidates randomly.

If the DOH service is unavailable, produces errors (HTTP or DNS), or the application otherwise fails to obtain an answer from it, the application MAY (but is not required to) fall back to using another configured DOH server, or to using "normal" DNS.

Likewise, hosts that do not match any configured bloom filter SHOULD be sent to a randomly selected DOH server that is available.

The means of discovering a DOH digest for a given DOH server is out of scope for this document, but generally it will be pre-arranged between the application and the DOH server.

The nature of this arrangement is highly dependent upon the application and its desired properties. That said, a number of requirements are placed upon this arrangement.

- o The digest MUST be conveyed in a manner that is secure and authenticated; e.g., TLS with appropriate certificate checks. Clients MUST enforce this.
- o The application MUST consider the DOH service as meeting whatever criteria it deems fit for configuring a "catch-all" DOH service (e.g., in terms of privacy, service availability, etc.), since false positives might be sent to the service, and hosts not matched by any configured bloom filter might be sent to it.

- o The digest MUST be updated on a periodic basis; e.g., once a day. Clients SHOULD NOT use stale digests.

3.2. The DOH Digest Format

TBD - likely just a bloom filter.

3.3. Hostname Normalisation

TBD

4. Security Considerations

Because a DOH digest allows a DOH server to claim traffic from an arbitrary hostname, applications need to take extreme care in selecting the DOH servers they will be accepted from, as well as assuring that their integrity and authentication have not been compromised.

Applications might mitigate this by monitoring DOH servers for such abuse and terminating their ability to use DOH digests when it is found.

TBD - more advanced mitigations

A hostname is effectively captured by a DOH server until the digest that reflects any change in its status is updated in the application. This delay should not result in any loss of functionality, since the "old" configuration will still direct requests to a functional DOH server.

5. IANA Considerations

This document currently has no IANA actions, but may grow some as the document progresses.

6. References

6.1. Normative References

- [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>>.
- [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>>.

6.2. Informative References

[fragile] Kashaf, A., Zarate, C., Wang, H., Agarwal, Y., and V. Sekar, "Oh, What a Fragile Web We Weave: Third-party Service Dependencies In Modern Webservices and Implications", June 2018, <<https://arxiv.org/pdf/1806.08420.pdf>>.

[I-D.ietf-doh-dns-over-https] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", [draft-ietf-doh-dns-over-https-12](#) (work in progress), June 2018.

[I-D.ietf-httpbis-http2-secondary-certs] Bishop, M., Sullivan, N., and M. Thomson, "Secondary Certificate Authentication in HTTP/2", [draft-ietf-httpbis-http2-secondary-certs-02](#) (work in progress), June 2018.

Author's Address

Mark Nottingham

Email: mnot@mnot.net

URI: <https://www.mnot.net/>