

HTTP Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 20, 2018

B. Schwartz
Google
M. Bishop
Akamai Technologies
January 16, 2018

Finding HTTP Alternative Services via the Domain Name Service
draft-schwartz-httpbis-dns-alt-svc-01

Abstract

The HTTP Alternative Services (Alt-Svc) mechanism allows an HTTP origin to be served from multiple network endpoints, and over multiple protocols. However, the client must first contact the origin server, in order to learn of the alternative services. This draft proposes a straightforward mapping of Alt-Svc into DNS, allowing clients to learn of these services before their first contact with the origin. This arrangement offers potential benefits to both performance and privacy.

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 July 20, 2018.

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	3
2.	The ALTSVC record type	3
2.1.	Comparison with alternatives	4
2.1.1.	Differences from the SRV RRTYPE	4
2.1.2.	Differences from the TXT RRTYPE	4
3.	Differences from Alt-Svc as transmitted over HTTP	5
3.1.	Omitting Max Age	5
3.2.	Interaction with other standards	5
3.3.	Granularity and lifetime control	5
4.	Client behaviors	6
4.1.	Cache interaction	6
4.2.	Optimizing for performance	6
4.3.	Optimizing for privacy	7
5.	Security Considerations	7
6.	IANA Considerations	8
7.	References	8
7.1.	Normative References	8
7.2.	Informative References	9
	Authors' Addresses	9

[1.](#) Introduction

The HTTP Alternative Services standard [[AltSvc](#)] defines

- o an extensible data model for describing alternative network endpoints that are authoritative for an origin
- o the "Alt-Svc Field Value", a text format for representing this information
- o standards for sending information in this format from a server to a client over HTTP/1.1 and HTTP/2.

Together, these components provide a toolkit that has proven useful and effective for informing a client of alternative services for an origin. However, making use of an alternative service requires contacting the origin server first. This creates an obvious performance cost: users wait for a full HTTP connection initiation (multiple roundtrips) before learning of an alternative service that is preferred by the origin. The first connection also publicly

reveals the user's intended destination to all entities along the network path.

This draft proposes a straightforward mechanism to distribute the Alt-Svc Field Value, in its standard text format, through the DNS. If a client receives this information during DNS resolution, it can skip the initial connection and proceed directly to an alternative service.

1.1. Terminology

For consistency with [\[AltSvc\]](#), we adopt the following definitions

- o An "origin" is an information source as in [\[RFC6454\]](#).
- o The "origin server" is the server that the client would reach when accessing the origin in the absence of Alt-Svc.
- o An "alternative service" is a different server that can serve the origin.

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. The ALTSVC record type

The ALTSVC DNS resource record (RR) type (RRTYPE ???) is used to associate an Alternative Service Field Value with an origin. Abstractly, the origin consists of a scheme (typically "https"), a host name, and a port (typically "443").

In the case of the ALTSVC RR, the origin is represented by prefixing the scheme and port with "_", then concatenating them with the host, resulting in a domain name like "_https._443.www.example.com."

The RDATA portion of an ALTSVC resource record contains an Alt-Svc Field Value, exactly as defined in Section 4 of [\[AltSvc\]](#).

For example, if the operator of https://www.example.com intends to include an HTTP response header like

```
Alt-Svc: h2=":8000"; ma=60
```

They would also publish an ALTSVC DNS record like


```
_https._443.www.example.com. 60S IN ALTSVC "h2=\":8000\""
```

This data type can be represented as an Unknown RR as described in [\[RFC3597\]](#):

```
_https._443.www.example.com. 60S IN TYPE??? \# 10  
68323D223A3830303022
```

This construction is intended to be extensible in two ways. First, any extensions that are made to the Alt-Svc format for transmission over HTTPS are also applicable here, unless expressly mentioned otherwise. Second, including the scheme in the DNS name allows for ALTSVC to serve schemes other than HTTPS, such as HTTP with Opportunistic Security [\[RFC8164\]](#) and any future schemes for which Alt-Svc may be defined.

[2.1.](#) Comparison with alternatives

The ALTSVC record type closely resembles some existing record types.

[2.1.1.](#) Differences from the SRV RRTYPE

An SRV record can perform a similar function to the ALTSVC record, informing a client to look in a different location for a service. However, there are several differences:

- o SRV records are typically mandatory, whereas clients will always continue to function correctly without making use of Alt-Svc.
- o SRV records cannot instruct the client to switch or upgrade protocols, whereas Alt-Svc can signal such an upgrade (e.g. to HTTP/2).
- o SRV records are not extensible, whereas Alt-Svc can be extended with new parameters. For example, this is what allows the privacy improvements related to SNI selection in [\[AltSvcSNI\]](#).
- o Using SRV records would not allow a client to skip processing of the Alt-Svc information in a subsequent connection, so it does not confer a performance advantage.

[2.1.2.](#) Differences from the TXT RRTYPE

The ALTSVC record uses an identical format to a TXT record, and could be implemented as such. However, we define a new record type for clarity, and to respect the use of TXT for human-readable notes as recommended in [\[RFC5507\]](#).

3. Differences from Alt-Svc as transmitted over HTTP

Publishing an ALTSVC record in DNS is intended to be equivalent to transmitting this field value over HTTP, and receiving an ALTSVC record is intended to be equivalent to receiving this field value over HTTP. However, there are some small differences in the intended client and server behavior.

3.1. Omitting Max Age

When publishing an ALTSVC record in DNS, server operators **MUST** omit the "ma" parameter, which encodes the "max age" (i.e. expiration time) of an Alt-Svc Field Value. Instead, server operators **SHOULD** encode the expiration time in the DNS TTL, and **MUST NOT** set a TTL longer than the intended "max age".

Server operators **MAY** publish multiple ALTSVC records as an RRSET, with semantics equivalent to other mechanisms of providing multiple Alt-Svc values to the client. When publishing an RRSET with multiple ALTSVC records, the server operator **MUST** set the overall TTL to the minimum of the "max age" values (following [Section 5.2 of \[RFC2181\]](#)).

When receiving an ALTSVC record, clients **MAY** synthesize a new "ma" parameter from the DNS TTL, in order to interoperate with Alt-Svc processing subsystems.

3.2. Interaction with other standards

The purpose of this standard is to reduce connection latency and improve user privacy. Server operators implementing this standard **SHOULD** also implement TLS 1.3 [[I-D.ietf-tls-tls13](#)] and OCSP Stapling [[RFC6066](#)], both of which confer substantial performance and privacy benefits when used in combination with ALTSVC records.

To realize the greatest privacy benefits, this proposal is intended for use with a privacy-preserving DNS transport (like DNS over TLS [[RFC7858](#)] or DNS over HTTPS [[DOH](#)]), and with the "SNI" Alt-Svc Parameter [[AltSvcSNI](#)]. However, performance improvements, and some modest privacy improvements, are possible without the use of those standards.

3.3. Granularity and lifetime control

Sending Alt-Svc over HTTP allows the server to tailor the Alt-Svc Field Value specifically to the client. When using an ALTSVC DNS record, groups of clients will necessarily receive the same Alt-Svc Field Value. Therefore, this standard is not suitable for servers that require single-client granularity in Alt-Svc. Server operators

that want to serve different Alt-Svc Field Values to different geographic or network regions SHOULD configure their authoritative DNS server to respect the EDNS0 Client Subnet extension [[RFC7871](#)].

Some DNS caching systems incorrectly extend the lifetime of DNS records beyond the stated TTL. Server operators MUST NOT rely on ALTSVC records expiring on time, and MAY shorten the TTL to compensate.

4. Client behaviors

4.1. Cache interaction

If the client has an Alt-Svc cache, and a usable Alt-Svc value is present in that cache, then the client SHOULD NOT issue an ALTSVC DNS query. Instead, the client SHOULD proceed with alternative service connection as usual.

If the client has a cached Alt-Svc entry that is expiring, the client MAY perform an ALTSVC query to refresh the entry.

4.2. Optimizing for performance

Clients that are optimizing for performance (i.e. minimum connection setup time) SHOULD implement the following connection sequence:

1. Issue address (AAAA and/or A) queries, immediately followed by the ALTSVC query.
2. If an ALTSVC response is received first, proceed with alternative service connection and ignore the address responses if they are no longer relevant.
3. Otherwise, initiate connection to the origin server.
4. As soon as an Alt-Svc field value is received, through the DNS or over HTTP, proceed with alternative service connection. Do not abort this connection if an Alt-Svc field value is received from the other source later.

If the ALTSVC and address queries return approximately simultaneously, this process typically saves three roundtrips on a fresh connection that uses Alt-Svc: one each for TCP, TLS 1.3, and HTTP. (On subsequent connections, the Alt-Svc information is expected to be cached, so this procedure does not apply.)

If a client can cache Alt-Svc entries that were received over both HTTP and DNS, the client MAY prefer entries that were received over

HTTP. These records may be more narrowly targeted for the specific client.

As an additional optimization, when choosing among multiple Alt-Svc values, clients MAY prefer those that will not require an address query, either because the corresponding address record is already in cache or because the host is an IP address.

Note that this procedure does not rely on recursive resolvers handling the ALTSVC record type correctly. If ALTSVC queries receive spurious NXDOMAIN responses, or even no response at all, connections will proceed as usual without any delay.

4.3. Optimizing for privacy

Clients that are optimizing for privacy SHOULD implement [[AltSvcSNI](#)] and DNS over a secure transport (e.g. [[RFC7858](#)] or [[DOH](#)]). Use of a secure transport is important not only for privacy protection, but also to ensure that queries for the new ALTSVC RRTYPE are handled correctly. Additionally, these clients SHOULD implement the following connection sequence:

1. Issue the ALTSVC DNS query first, immediately followed by the address queries.
2. Wait for the ALTSVC record response.
3. If the response is nonempty, proceed with alternative service connection and ignore the address query responses.
4. Otherwise, wait for the address queries and connect as usual.

Note that this process is also expected to be faster than Alt-Svc over HTTP in the case of HTTP Opportunistic Upgrade Probing ([Section 2 of \[RFC8164\]](#)).

5. Security Considerations

Alt-Svc Field Values are intended for distribution over untrusted channels, and clients are REQUIRED to verify that the alternative service is authoritative for the origin (Section 2.1 of [[AltSvc](#)]). Therefore, DNSSEC signing and validation are OPTIONAL for publishing and using ALTSVC records.

6. IANA Considerations

This draft requires assignment of a new DNS RRTYPE value.

7. References

7.1. Normative References

- [AltSvc] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", [RFC 7838](#), DOI 10.17487/RFC7838, April 2016, <<https://www.rfc-editor.org/info/rfc7838>>.
- [AltSvcSNI] Bishop, M., "The "SNI" Alt-Svc Parameter", [draft-bishop-httpbis-sni-altsvc-01](#) (work in progress), January 2018.
- [DOH] Hoffman, P. and P. McManus, "DNS Queries over HTTPS", [draft-ietf-doh-dns-over-https-02](#) (work in progress), November 2017.
- [I-D.ietf-tls-tls13] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", [draft-ietf-tls-tls13-23](#) (work in progress), January 2018.
- [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>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", [RFC 2181](#), DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.
- [RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", [RFC 3597](#), DOI 10.17487/RFC3597, September 2003, <<https://www.rfc-editor.org/info/rfc3597>>.
- [RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", [RFC 6066](#), DOI 10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/info/rfc6066>>.
- [RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/info/rfc6454>>.

- [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>>.

[7.2.](#) Informative References

- [RFC5507] IAB, Faltstrom, P., Ed., Austein, R., Ed., and P. Koch, Ed., "Design Choices When Expanding the DNS", [RFC 5507](#), DOI 10.17487/RFC5507, April 2009, <<https://www.rfc-editor.org/info/rfc5507>>.
- [RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", [RFC 7871](#), DOI 10.17487/RFC7871, May 2016, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC8164] Nottingham, M. and M. Thomson, "Opportunistic Security for HTTP/2", [RFC 8164](#), DOI 10.17487/RFC8164, May 2017, <<https://www.rfc-editor.org/info/rfc8164>>.

Authors' Addresses

Ben Schwartz
Google

Email: bemasc@google.com

Mike Bishop
Akamai Technologies

Email: mbishop@evequefou.be

