

DNSOP Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 6, 2021

B. Schwartz  
Google  
M. Bishop  
E. Nygren  
Akamai Technologies  
November 2, 2020

**Service binding and parameter specification via the DNS (DNS SVCB and  
HTTPS RRs)  
draft-ietf-dnsop-svcb-https-02**

**Abstract**

This document specifies the "SVCB" and "HTTPS" DNS resource record (RR) types to facilitate the lookup of information needed to make connections to network services, such as for HTTPS origins. SVCB records allow a service to be provided from multiple alternative endpoints, each with associated parameters (such as transport protocol configuration and keys for encrypting the TLS ClientHello). They also enable aliasing of apex domains, which is not possible with CNAME. The HTTPS RR is a variation of SVCB for HTTPS and HTTP origins. By providing more information to the client before it attempts to establish a connection, these records offer potential benefits to both performance and privacy.

TO BE REMOVED: This document is being collaborated on in Github at: <https://github.com/MikeBishop/dns-alt-svc> [1]. The most recent working version of the document, open issues, etc. should all be available there. The authors (gratefully) accept pull requests.

**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 May 6, 2021.

## Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">4</a>
<a href="#">1.1.</a>	<a href="#">Goals of the SVCB RR</a>	<a href="#">5</a>
<a href="#">1.2.</a>	<a href="#">Overview of the SVCB RR</a>	<a href="#">5</a>
<a href="#">1.3.</a>	<a href="#">Parameter for Encrypted ClientHello</a>	<a href="#">6</a>
<a href="#">1.4.</a>	<a href="#">Terminology</a>	<a href="#">7</a>
<a href="#">2.</a>	<a href="#">The SVCB record type</a>	<a href="#">7</a>
<a href="#">2.1.</a>	<a href="#">Zone file presentation format</a>	<a href="#">8</a>
<a href="#">2.2.</a>	<a href="#">RDATA wire format</a>	<a href="#">9</a>
<a href="#">2.3.</a>	<a href="#">SVCB query names</a>	<a href="#">10</a>
<a href="#">2.4.</a>	<a href="#">Interpretation</a>	<a href="#">10</a>
<a href="#">2.4.1.</a>	<a href="#">SvcPriority</a>	<a href="#">11</a>
<a href="#">2.4.2.</a>	<a href="#">AliasMode</a>	<a href="#">11</a>
<a href="#">2.4.3.</a>	<a href="#">ServiceMode</a>	<a href="#">12</a>
<a href="#">2.5.</a>	<a href="#">Special handling of "." in TargetName</a>	<a href="#">12</a>
<a href="#">2.5.1.</a>	<a href="#">AliasMode</a>	<a href="#">13</a>
<a href="#">2.5.2.</a>	<a href="#">ServiceMode</a>	<a href="#">13</a>
<a href="#">3.</a>	<a href="#">Client behavior</a>	<a href="#">13</a>
<a href="#">3.1.</a>	<a href="#">Handling resolution failures</a>	<a href="#">14</a>
<a href="#">3.2.</a>	<a href="#">Clients using a Proxy</a>	<a href="#">15</a>
<a href="#">4.</a>	<a href="#">DNS Server Behavior</a>	<a href="#">15</a>
<a href="#">4.1.</a>	<a href="#">Authoritative servers</a>	<a href="#">15</a>
<a href="#">4.2.</a>	<a href="#">Recursive resolvers</a>	<a href="#">16</a>
<a href="#">4.3.</a>	<a href="#">General requirements</a>	<a href="#">16</a>
<a href="#">5.</a>	<a href="#">Performance optimizations</a>	<a href="#">17</a>
<a href="#">5.1.</a>	<a href="#">Optimistic pre-connection and connection reuse</a>	<a href="#">17</a>
<a href="#">5.2.</a>	<a href="#">Generating and using incomplete responses</a>	<a href="#">18</a>
<a href="#">6.</a>	<a href="#">Initial SvcParamKeys</a>	<a href="#">18</a>
<a href="#">6.1.</a>	<a href="#">"alpn" and "no-default-alpn"</a>	<a href="#">18</a>
<a href="#">6.2.</a>	<a href="#">"port"</a>	<a href="#">20</a>
<a href="#">6.3.</a>	<a href="#">"echconfig"</a>	<a href="#">20</a>
<a href="#">6.4.</a>	<a href="#">"ipv4hint" and "ipv6hint"</a>	<a href="#">21</a>



<a href="#">7.</a>	<a href="#">ServiceMode RR compatibility and mandatory keys . . . . .</a>	<a href="#">22</a>
<a href="#">8.</a>	<a href="#">Using SVCB with HTTPS and HTTP . . . . .</a>	<a href="#">22</a>
<a href="#">8.1.</a>	<a href="#">Query names for HTTPS RRs . . . . .</a>	<a href="#">23</a>
<a href="#">8.2.</a>	<a href="#">Relationship to Alt-Svc . . . . .</a>	<a href="#">24</a>
<a href="#">8.2.1.</a>	<a href="#">ALPN usage . . . . .</a>	<a href="#">24</a>
<a href="#">8.2.2.</a>	<a href="#">Untrusted channel . . . . .</a>	<a href="#">24</a>
<a href="#">8.2.3.</a>	<a href="#">Cache lifetime . . . . .</a>	<a href="#">24</a>
<a href="#">8.2.4.</a>	<a href="#">Granularity . . . . .</a>	<a href="#">25</a>
<a href="#">8.3.</a>	<a href="#">Interaction with Alt-Svc . . . . .</a>	<a href="#">25</a>
<a href="#">8.4.</a>	<a href="#">Requiring Server Name Indication . . . . .</a>	<a href="#">25</a>
<a href="#">8.5.</a>	<a href="#">HTTP Strict Transport Security . . . . .</a>	<a href="#">25</a>
<a href="#">8.6.</a>	<a href="#">HTTP-based protocols . . . . .</a>	<a href="#">26</a>
<a href="#">9.</a>	<a href="#">SVCB/HTTPS RR parameter for ECH configuration . . . . .</a>	<a href="#">26</a>
<a href="#">9.1.</a>	<a href="#">Client behavior . . . . .</a>	<a href="#">27</a>
<a href="#">9.2.</a>	<a href="#">Deployment considerations . . . . .</a>	<a href="#">27</a>
<a href="#">10.</a>	<a href="#">Zone Structures . . . . .</a>	<a href="#">27</a>
<a href="#">10.1.</a>	<a href="#">Structuring zones for flexibility . . . . .</a>	<a href="#">27</a>
<a href="#">10.2.</a>	<a href="#">Structuring zones for performance . . . . .</a>	<a href="#">28</a>
<a href="#">10.3.</a>	<a href="#">Examples . . . . .</a>	<a href="#">28</a>
<a href="#">10.3.1.</a>	<a href="#">Protocol enhancements . . . . .</a>	<a href="#">28</a>
<a href="#">10.3.2.</a>	<a href="#">Apex aliasing . . . . .</a>	<a href="#">28</a>
<a href="#">10.3.3.</a>	<a href="#">Parameter binding . . . . .</a>	<a href="#">29</a>
<a href="#">10.3.4.</a>	<a href="#">Multi-CDN . . . . .</a>	<a href="#">29</a>
<a href="#">10.3.5.</a>	<a href="#">Non-HTTPS uses . . . . .</a>	<a href="#">31</a>
<a href="#">11.</a>	<a href="#">Interaction with other standards . . . . .</a>	<a href="#">32</a>
<a href="#">12.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">32</a>
<a href="#">13.</a>	<a href="#">Privacy Considerations . . . . .</a>	<a href="#">33</a>
<a href="#">14.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">33</a>
<a href="#">14.1.</a>	<a href="#">SVCB RRTYPE . . . . .</a>	<a href="#">33</a>
<a href="#">14.2.</a>	<a href="#">HTTPS RRTYPE . . . . .</a>	<a href="#">33</a>
<a href="#">14.3.</a>	<a href="#">New registry for Service Parameters . . . . .</a>	<a href="#">34</a>
<a href="#">14.3.1.</a>	<a href="#">Procedure . . . . .</a>	<a href="#">34</a>
<a href="#">14.3.2.</a>	<a href="#">Initial contents . . . . .</a>	<a href="#">34</a>
<a href="#">14.4.</a>	<a href="#">Registry updates . . . . .</a>	<a href="#">35</a>
<a href="#">15.</a>	<a href="#">Acknowledgments and Related Proposals . . . . .</a>	<a href="#">36</a>
<a href="#">16.</a>	<a href="#">References . . . . .</a>	<a href="#">36</a>
<a href="#">16.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">36</a>
<a href="#">16.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">39</a>
<a href="#">16.3.</a>	<a href="#">URIs . . . . .</a>	<a href="#">40</a>
<a href="#">Appendix A.</a>	<a href="#">Decoding text in zone files . . . . .</a>	<a href="#">40</a>
<a href="#">A.1.</a>	<a href="#">Decoding a value-list . . . . .</a>	<a href="#">40</a>
<a href="#">Appendix B.</a>	<a href="#">Comparison with alternatives . . . . .</a>	<a href="#">41</a>
<a href="#">B.1.</a>	<a href="#">Differences from the SRV RR type . . . . .</a>	<a href="#">41</a>
<a href="#">B.2.</a>	<a href="#">Differences from the proposed HTTP record . . . . .</a>	<a href="#">41</a>
<a href="#">B.3.</a>	<a href="#">Differences from the proposed ANAME record . . . . .</a>	<a href="#">42</a>
<a href="#">B.4.</a>	<a href="#">Comparison with separate RR types for AliasMode and ServiceMode . . . . .</a>	<a href="#">42</a>
<a href="#">Appendix C.</a>	<a href="#">Change history . . . . .</a>	<a href="#">42</a>



Authors' Addresses . . . . .	<a href="#">45</a>
------------------------------	--------------------

## **[1.](#) Introduction**

The SVCB ("Service Binding") and HTTPS RRs provide clients with complete instructions for access to a service. This information enables improved performance and privacy by avoiding transient connections to a sub-optimal default server, negotiating a preferred protocol, and providing relevant public keys.

For example, when clients need to make a connection to fetch resources associated with an HTTPS URI, they currently resolve only A and/or AAAA records for the origin hostname. This is adequate for services that use basic HTTPS (fixed port, no QUIC, no [\[ECH\]](#)). Going beyond basic HTTPS confers privacy, performance, and operational advantages, but it requires the client to learn additional information, and it is highly desirable to minimize the number of round-trips and lookups required to learn this additional information.

The SVCB and HTTPS RRs also help when the operator of a service wishes to delegate operational control to one or more other domains, e.g. delegating the origin "https://example.com" to a service operator endpoint at "svc.example.net". While this case can sometimes be handled by a CNAME, that does not cover all use-cases. CNAME is also inadequate when the service operator needs to provide a bound collection of consistent configuration parameters through the DNS (such as network location, protocol, and keying information).

This document first describes the SVCB RR as a general-purpose resource record that can be applied directly and efficiently to a wide range of services ([Section 2](#)). The HTTPS RR is then defined as a special case of SVCB that improves efficiency and convenience for use with HTTPS ([Section 8](#)) by avoiding the need for an Attrleaf label [[Attrleaf](#)] ([Section 8.1](#)). Other protocols with similar needs may follow the pattern of HTTPS and assign their own SVCB-compatible RR types.

All behaviors described as applying to the SVCB RR also apply to the HTTPS RR unless explicitly stated otherwise. [Section 8](#) describes additional behaviors specific to the HTTPS RR. Apart from [Section 8](#) and introductory examples, much of this document refers only to the SVCB RR, but those references should be taken to apply to SVCB, HTTPS, and any future SVCB-compatible RR types.

The SVCB RR has two modes: 1) "AliasMode" simply delegates operational control for a resource; 2) "ServiceMode" binds together



configuration information for a service endpoint. ServiceMode provides additional key=value parameters within each RDATA set.

### **1.1. Goals of the SVCB RR**

The goal of the SVCB RR is to allow clients to resolve a single additional DNS RR in a way that:

- o Provides alternative endpoints that are authoritative for the service, along with parameters associated with each of these endpoints.
- o Does not assume that all alternative endpoints have the same parameters or capabilities, or are even operated by the same entity. This is important as DNS does not provide any way to tie together multiple RRs for the same name. For example, if `www.example.com` is a CNAME alias that switches between one of three CDNs or hosting environments, successive queries for that name may return records that correspond to different environments.
- o Enables CNAME-like functionality at a zone apex (such as "example.com") for participating protocols, and generally enables delegation of operational authority for an origin within the DNS to an alternate name.

Additional goals specific to HTTPS RRs and the HTTPS use-case include:

- o Connect directly to HTTP3 (QUIC transport) alternative endpoints [[HTTP3](#)]
- o Obtain the Encrypted ClientHello [[ECH](#)] keys associated with an alternative endpoint
- o Support non-default TCP and UDP ports
- o Enable SRV-like benefits (e.g. apex delegation, as mentioned above) for HTTP(S), where SRV [[SRV](#)] has not been widely adopted
- o Provide an HSTS-like indication [[HSTS](#)] signaling that the HTTPS scheme should be used instead of HTTP for this request (see [Section 8.5](#)).

### **1.2. Overview of the SVCB RR**

This subsection briefly describes the SVCB RR in a non-normative manner. (As mentioned above, this all applies equally to the HTTPS RR which shares the same encoding, format, and high-level semantics.)





The SVCB RR has two modes: AliasMode, which aliases a name to another name, and ServiceMode, which provides connection information bound to a service endpoint domain. Placing both forms in a single RR type allows clients to fetch the relevant information with a single query.

The SVCB RR has two mandatory fields and one optional. The fields are:

1. SvcPriority: The priority of this record (relative to others, with lower values preferred). A value of 0 indicates AliasMode. (Described in [Section 2.4.1](#).)
2. TargetName: The domain name of either the alias target (for AliasMode) or the alternative endpoint (for ServiceMode).
3. SvcParams (optional): A list of key=value pairs describing the alternative endpoint at TargetName (only used in ServiceMode and otherwise ignored). Described in [Section 2.1](#).

Cooperating DNS recursive resolvers will perform subsequent record resolution (for SVCB, A, and AAAA records) and return them in the Additional Section of the response. Clients either use responses included in the additional section returned by the recursive resolver or perform necessary SVCB, A, and AAAA record resolutions. DNS authoritative servers can attach in-bailiwick SVCB, A, AAAA, and CNAME records in the Additional Section to responses for a SVCB query.

In ServiceMode, the SvcParams of the SVCB RR provide an extensible data model for describing alternative endpoints that are authoritative for the origin, along with parameters associated with each of these alternative endpoints.

For the HTTPS use-case, the HTTPS RR enables many of the benefits of Alt-Svc [[AltSvc](#)] without waiting for a full HTTP connection initiation (multiple roundtrips) before learning of the preferred alternative, and without necessarily revealing the user's intended destination to all entities along the network path.

### **[1.3](#). Parameter for Encrypted ClientHello**

This document also defines a parameter for Encrypted ClientHello [[ECH](#)] keys. See [Section 9](#).



## 1.4. Terminology

Our terminology is based on the common case where the SVCB record is used to access a resource identified by a URI whose "authority" field contains a DNS hostname as the "host".

- o The "service" is the information source identified by the "authority" and "scheme" of the URI, capable of providing access to the resource. For HTTPS URIs, the "service" corresponds to an HTTPS "origin" [[RFC6454](#)].
- o The "service name" is the "host" portion of the authority.
- o The "authority endpoint" is the authority's hostname and a port number implied by the scheme or specified in the URI.
- o An "alternative endpoint" is a hostname, port number, and other associated instructions to the client on how to reach an instance of service.

Additional DNS terminology intends to be consistent with [[DNSTerm](#)].

SVCB is a contraction of "service binding". The SVCB RR, HTTPS RR, and future RR types that share SVCB's format and registry are collectively known as SVCB-compatible RR types.

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 SVCB record type

The SVCB DNS resource record (RR) type (RR type 64) is used to locate alternative endpoints for a service.

The algorithm for resolving SVCB records and associated address records is specified in [Section 3](#).

Other SVCB-compatible resource record types can also be defined as-needed. In particular, the HTTPS RR (RR type 65) provides special handling for the case of "https" origins as described in [Section 8](#).

SVCB RRs are extensible by a list of SvcParams, which are pairs consisting of a SvcParamKey and a SvcParamValue. Each SvcParamKey has a presentation name and a registered number. Values are in a format specific to the SvcParamKey. Their definition should specify



both their presentation format and wire encoding (e.g., domain names, binary data, or numeric values). The initial SvcParamKeys and formats are defined in [Section 6](#).

### **2.1. Zone file presentation format**

The presentation format of the record is:

```
Name TTL IN SVCB SvcPriority TargetName SvcParams
```

The SVCB record is defined specifically within the Internet ("IN") Class ([RFC1035](#)).

SvcPriority is a number in the range 0-65535, TargetName is a domain name, and the SvcParams are a whitespace-separated list, with each SvcParam consisting of a SvcParamKey=SvcParamValue pair or a standalone SvcParamKey. SvcParamKeys are subject to IANA control ([Section 14.3](#)).

Each SvcParamKey SHALL appear at most once in the SvcParams. In presentation format, SvcParamKeys are lower-case alphanumeric strings. Key names should contain 1-63 characters from the ranges "a"-"z", "0"-"9", and "-". In ABNF [RFC5234](#),

```
alpha-lc      = %x61-7A ; a-z
SvcParamKey   = 1*63(alpha-lc / DIGIT / "-")
SvcParam      = SvcParamKey ["=" SvcParamValue]
SvcParamValue = char-string
value         = *OCTET
```

Unless otherwise specified, the SvcParamValue is parsed using the character-string decoding algorithm (Appendix A), producing a "value". The "value" is then validated and converted into wire-format in a manner specific to each key.

When the "=" is omitted, the "value" is interpreted as empty.

Unrecognized keys are represented in presentation format as "keyNNNNN" where NNNNN is the numeric value of the key type without leading zeros. A SvcParam in this form SHALL be parsed as specified above, and the decoded "value" SHALL be used as its wire format encoding.

For some SvcParamKeys, the SvcParamValue corresponds to a list or set of items. Presentation formats for such keys SHOULD make use of the decoding algorithm in [Appendix A.1](#), producing a "value-list".



SvcParams in presentation format MAY appear in any order, but keys MUST NOT be repeated.

## 2.2. RDATA wire format

The RDATA for the SVCB RR consists of:

- o a 2 octet field for SvcPriority as an integer in network byte order.
- o the uncompressed, fully-qualified TargetName, represented as a sequence of length-prefixed labels as in [Section 3.1 of \[RFC1035\]](#).
- o the SvcParams, consuming the remainder of the record (so smaller than 65535 octets and constrained by the RDATA and DNS message sizes).

When the list of SvcParams is non-empty (ServiceMode), it contains a series of SvcParamKey=SvcParamValue pairs, represented as:

- o a 2 octet field containing the SvcParamKey as an integer in network byte order. (See [Section 14.3.2](#) for the defined values.)
- o a 2 octet field containing the length of the SvcParamValue as an integer between 0 and 65535 in network byte order (but constrained by the RDATA and DNS message sizes).
- o an octet string of this length whose contents are in a format determined by the SvcParamKey.

SvcParamKeys SHALL appear in increasing numeric order.

Clients MUST consider an RR malformed if:

- o the end of the RDATA occurs within a SvcParam.
- o SvcParamKeys are not in strictly increasing numeric order.
- o the SvcParamValue for an SvcParamKey does not have the expected format.

Note that the second condition implies that there are no duplicate SvcParamKeys.

If any RRs are malformed, the client MUST reject the entire RRSets and fall back to non-SVCB connection establishment.





### **2.3. SVCB query names**

When querying the SVCB RR, a service is translated into a QNAME by prepending the service name with a label indicating the scheme, prefixed with an underscore, resulting in a domain name like "\_examplescheme.api.example.com.". This follows the Attrleaf naming pattern [[Attrleaf](#)], so the scheme MUST be registered appropriately with IANA (see [Section 11](#)).

Protocol mapping documents MAY specify additional underscore-prefixed labels to be prepended. For schemes that specify a port (Section 3.2.3 of [[URI](#)]), one reasonable possibility is to prepend the indicated port number (or the default if no port number is specified). We term this behavior "Port Prefix Naming", and use it in the examples throughout this document.

See [Section 8.1](#) for the HTTPS RR behavior.

When a prior CNAME or SVCB record has aliased to a SVCB record, each RR shall be returned under its own owner name.

Note that none of these forms alter the origin or authority for validation purposes. For example, TLS clients MUST continue to validate TLS certificates for the original service name.

As an example, the owner of example.com could publish this record:

```
_8443._foo.api.example.com. 7200 IN SVCB 0 svc4.example.net.
```

to indicate that "foo://api.example.com:8443" is aliased to "svc4.example.net". The owner of example.net, in turn, could publish this record:

```
svc4.example.net. 7200 IN SVCB 3 svc4.example.net. (  
    alpn="bar" port="8004" echconfig="..." )
```

to indicate that these services are served on port number 8004, which supports the protocol "bar" and its associated transport in addition to the default transport protocol for "foo://".

(Parentheses are used to ignore a line break ([\[RFC1035\]](#)  
[Section 5.1](#))).)

### **2.4. Interpretation**



#### **2.4.1. SvcPriority**

When SvcPriority is 0 the SVCB record is in AliasMode ([Section 2.4.2](#)). Otherwise, it is in ServiceMode ([Section 2.4.3](#)).

Within a SVCB RSet, all RRs SHOULD have the same Mode. If an RSet contains a record in AliasMode, the recipient MUST ignore any ServiceMode records in the set.

RRSets are explicitly unordered collections, so the SvcPriority field is used to impose an ordering on SVCB RRs. SVCB RRs with a smaller SvcPriority value SHOULD be given preference over RRs with a larger SvcPriority value.

When receiving an RSet containing multiple SVCB records with the same SvcPriority value, clients SHOULD apply a random shuffle within a priority level to the records before using them, to ensure uniform load-balancing.

#### **2.4.2. AliasMode**

In AliasMode, the SVCB record aliases a service to a TargetName. SVCB RSets SHOULD only have a single resource record in AliasMode. If multiple are present, clients or recursive resolvers SHOULD pick one at random.

The primary purpose of AliasMode is to allow aliasing at the zone apex, where CNAME is not allowed. In AliasMode, the TargetName will be the name of a domain that resolves to SVCB (or other SVCB-compatible record such as HTTPS), AAAA, and/or A records. The TargetName SHOULD NOT be equal to the owner name, as this would result in a loop.

In AliasMode, records SHOULD NOT include any SvcParams, and recipients MUST ignore any SvcParams that are present.

For example, the operator of foo://example.com:8080 could point requests to a service operating at foosvc.example.net by publishing:

```
_8080._foo.example.com. 3600 IN SVCB 0 foosvc.example.net.
```

Using AliasMode maintains a separation of concerns: the owner of foosvc.example.net can add or remove ServiceMode SVCB records without requiring a corresponding change to example.com. Note that if foosvc.example.net promises to always publish a SVCB record, this AliasMode record can be replaced by a CNAME, which would likely improve performance.



AliasMode is especially useful for SVCB-compatible RR types that do not require an underscore prefix, such as the HTTPS RR type. For example, the operator of `https://example.com` could point requests to a server at `svc.example.net` by publishing this record at the zone apex:

```
example.com. 3600 IN HTTPS 0 svc.example.net.
```

Note that the SVCB record's owner name MAY be the canonical name of a CNAME record, and the TargetName MAY be the owner of a CNAME record. Clients and recursive resolvers MUST follow CNAMEs as normal.

To avoid unbounded alias chains, clients and recursive resolvers MUST impose a limit on the total number of SVCB aliases they will follow for each resolution request. This limit MUST NOT be zero, i.e. implementations MUST be able to follow at least one AliasMode record. The exact value of this limit is left to implementations.

For compatibility and performance, zone owners SHOULD NOT configure their zones to require following multiple AliasMode records.

As legacy clients will not know to use this record, service operators will likely need to retain fallback AAAA and A records alongside this SVCB record, although in a common case the target of the SVCB record might offer better performance, and therefore would be preferable for clients implementing this specification to use.

AliasMode records only apply to queries for the specific RR type. For example, a SVCB record cannot alias to an HTTPS record, nor vice-versa.

#### **2.4.3. ServiceMode**

In ServiceMode, the TargetName and SvcParams within each resource record associate an alternative endpoint for the service with its connection parameters.

Each protocol scheme that uses SVCB MUST define a protocol mapping that explains how SvcParams are applied for connections of that scheme. Unless specified otherwise by the protocol mapping, clients MUST ignore any SvcParam that they do not recognize.

#### **2.5. Special handling of "." in TargetName**

If TargetName has the value "." (represented in the wire format as a zero-length label), special rules apply.



### [2.5.1.](#) AliasMode

For AliasMode SVCB RRs, a TargetName of "." indicates that the service is not available or does not exist. This indication is advisory: clients encountering this indication MAY ignore it and attempt to connect without the use of SVCB.

### [2.5.2.](#) ServiceMode

For ServiceMode SVCB RRs, if TargetName has the value ".", then the owner name of this record MUST be used as the effective TargetName.

For example, in the following example "svc2.example.net" is the effective TargetName:

```
example.com.      7200  IN HTTPS 0  svc.example.net.
svc.example.net.  7200  IN CNAME  svc2.example.net.
svc2.example.net. 7200  IN HTTPS 1  . port=8002 echconfig="..."
svc2.example.net. 300   IN A      192.0.2.2
svc2.example.net. 300   IN AAAA   2001:db8::2
```

## [3.](#) Client behavior

A SVCB-aware client selects an endpoint for a service using the following procedure:

1. Let \$ADDR\_QNAME be the service name. Let \$SVCB\_QNAME be the service name plus appropriate prefixes for the scheme (see [Section 2.3](#)).
2. In parallel, issue AAAA/A queries for \$ADDR\_QNAME and a SVCB query for \$SVCB\_QNAME. The answers for these may or may not include CNAME pointers before reaching one or more of these records.
3. If an AliasMode SVCB record is returned for \$SVCB\_QNAME, clients MUST set \$ADDR\_QNAME and \$SVCB\_QNAME to its TargetName (without additional prefixes) and loop back to step 2, subject to chain length limits and loop detection heuristics (see [Section 3.1](#)).
4. If one or more "compatible" ([Section 7](#)) ServiceMode records are returned for \$SVCB\_QNAME, clients SHOULD select the highest-priority compatible record. This record's TargetName and SvcParams represent the preferred endpoint. If connection to this endpoint fails, the client SHOULD try to connect using values from the next-highest-priority compatible record, etc. If all attempts fail, clients SHOULD go to step 5 (except as noted in [Section 9.1](#)).





5. At this point there are no usable ServiceMode records, because

- \* there were no SVCB records found for \$SVCB\_QNAME, OR
- \* all records found were incompatible with this client, OR
- \* all connection attempts using ServiceMode records failed.

Accordingly, clients SHALL connect to the endpoint consisting of \$ADDR\_QNAME, the authority endpoint's port number, and no SvcParams.

If all of the above connection attempts fail, clients MAY connect to the authority endpoint (except as noted in [Section 3.1](#) and [Section 9.1](#)).

This procedure does not rely on any recursive or authoritative DNS server to comply with this specification or have any awareness of SVCB.

When selecting between AAAA and A records to use, clients may use an approach such as Happy Eyeballs [[HappyEyeballsV2](#)].

Some important optimizations are discussed in [Section 5](#) to avoid additional latency in comparison to ordinary AAAA/A lookups.

### **[3.1](#). Handling resolution failures**

If a SVCB query results in a SERVFAIL error, transport error, or timeout, and DNS exchanges between the client and the recursive resolver are cryptographically protected (e.g. using TLS [[DoT](#)] or HTTPS [[DoH](#)]), the client SHOULD NOT fall back to \$ADDR\_QNAME (step 5 above) or the authority endpoint. Otherwise, an active attacker could mount a downgrade attack by denying the user access to the SvcParams.

A SERVFAIL error can occur if the domain is DNSSEC-signed, the recursive resolver is DNSSEC-validating, and the attacker is between the recursive resolver and the authoritative DNS server. A transport error or timeout can occur if an active attacker between the client and the recursive resolver is selectively dropping SVCB queries or responses, based on their size or other observable patterns.

Similarly, if the client enforces DNSSEC validation on A/AAAA responses, it SHOULD terminate the connection if a SVCB response fails to validate.



If the client is unable to complete SVCB resolution due to its chain length limit, the client SHOULD fall back to the authority endpoint, as if the origin's SVCB record did not exist.

### **3.2. Clients using a Proxy**

Clients using a domain-oriented transport proxy like HTTP CONNECT ([RFC7231] Section 4.3.6) or SOCKS5 ([RFC1928]) have the option to use named destinations, in which case the client does not perform any A or AAAA queries for destination domains. If the client is using named destinations with a proxy that does not provide SVCB query capability (e.g. through an affiliated DNS resolver), the client would have to perform SVCB queries through a separate resolver. This might disclose the client's destinations to an additional party, creating privacy concerns. If these concerns apply, the client SHOULD disable SVCB resolution.

If the client does use SVCB and named destinations, the client SHOULD follow the standard SVCB resolution process, selecting the smallest-SvcPriority option that is compatible with the client and the proxy. The client SHOULD provide the final TargetName and port to the proxy, which will perform any required A and AAAA lookups.

Providing the proxy with the final TargetName has several benefits:

- o It allows the client to use the SvcParams, if present, which is only usable with a specific TargetName. The SvcParams may include information that enhances performance (e.g. alpn) and privacy (e.g. echconfig).
- o It allows the service to delegate the apex domain.
- o It allows the proxy to select between IPv4 and IPv6 addresses for the server according to its configuration, and receive addresses based on its network geolocation.

## **4. DNS Server Behavior**

### **4.1. Authoritative servers**

When replying to a SVCB query, authoritative DNS servers SHOULD return A, AAAA, and SVCB records in the Additional Section for any in-bailiwick TargetNames. If the zone is signed, the server SHOULD also include positive or negative DNSSEC responses for these records in the Additional section.



## 4.2. Recursive resolvers

Recursive resolvers that are aware of SVCB SHOULD help the client to execute the procedure in [Section 3](#) with minimum overall latency, by incorporating additional useful information into the response. For the initial SVCB record query, this is just the normal response construction process (i.e. unknown RR type resolution under [\[RFC3597\]](#)). For followup resolutions performed during this procedure, we define incorporation as adding all useful RRs from the response to the Additional section without altering the response code.

Upon receiving a SVCB query, recursive resolvers SHOULD start with the standard resolution procedure, and then follow this procedure to construct the full response to the stub resolver:

1. Incorporate the results of SVCB resolution. If the chain length limit has been reached, terminate successfully (i.e. a NOERROR response).
2. If any of the resolved SVCB records are in AliasMode, choose one of them at random, and resolve SVCB, A, and AAAA records for its TargetName.
  - \* If any SVCB records are resolved, go to step 1.
  - \* Otherwise, incorporate the results of A and AAAA resolution, and terminate.
3. All the resolved SVCB records are in ServiceMode. Resolve A and AAAA queries for each TargetName (or for the owner name if TargetName is "."), incorporate all the results, and terminate.

In this procedure, "resolve" means the resolver's ordinary recursive resolution procedure, as if processing a query for that RRSets. This includes following any aliases that the resolver would ordinarily follow (e.g. CNAME, DNAME [\[DNAME\]](#)).

See [Section 2.4.2](#) for additional safeguards for recursive resolvers to implement to mitigate loops.

See [Section 5.2](#) for possible optimizations of this procedure.

## 4.3. General requirements

Recursive resolvers SHOULD treat the SvcParams portion of the SVCB RR as opaque and SHOULD NOT try to alter their behavior based on its contents.



When responding to a query that includes the DNSSEC OK bit ([RFC3225]), DNSSEC-capable recursive and authoritative DNS servers MUST accompany each RRSSet in the Additional section with the same DNSSEC-related records that they would send when providing that RRSSet as an Answer (e.g. RRSIG, NSEC, NSEC3).

## 5. Performance optimizations

For optimal performance (i.e. minimum connection setup time), clients SHOULD issue address (AAAA and/or A) and SVCB queries simultaneously, and SHOULD implement a client-side DNS cache. Responses in the Additional section of a SVCB response SHOULD be placed in cache before performing any followup queries. With these optimizations in place, and conforming DNS servers, using SVCB does not add network latency to connection setup.

### 5.1. Optimistic pre-connection and connection reuse

If an address response arrives before the corresponding SVCB response, the client MAY initiate a connection as if the SVCB query returned NODATA, but MUST NOT transmit any information that could be altered by the SVCB response until it arrives. For example, a TLS ClientHello can be altered by the "echconfig" value of a SVCB response (Section 6.3). Clients implementing this optimization SHOULD wait for 50 milliseconds before starting optimistic pre-connection, as per the guidance in [HappyEyeballsV2].

An SVCB record is consistent with a connection if the client would attempt an equivalent connection when making use of that record. If a SVCB record is consistent with an active or in-progress connection C, the client MAY prefer that record and use C as its connection. For example, suppose the client receives this SVCB RRSSet for a protocol that uses TLS over TCP:

```
_1234._bar.example.com. 300 IN SVCB 1 svc1.example.net. (
    echconfig="111..." ipv6hint=2001:db8::1 port=1234 )
                        SVCB 2 svc2.example.net. (
    echconfig="222..." ipv6hint=2001:db8::2 port=1234 )
```

If the client has an in-progress TCP connection to "[2001:db8::2]:1234", it MAY proceed with TLS on that connection using "echconfig="222...", even though the other record in the RRSSet has higher priority.

If none of the SVCB records are consistent with any active or in-progress connection, clients must proceed as described in Step 3 of the procedure in Section 3.





## 5.2. Generating and using incomplete responses

When following the procedure in [Section 4.2](#), recursive resolvers MAY terminate the procedure early and produce a reply that omits some of the associated RRSets. This is REQUIRED when the chain length limit is reached ([Section 4.2](#) step 1), but might also be appropriate when the maximum response size is reached, or when responding before fully chasing dependencies would improve performance. When omitting certain RRSets, recursive resolvers SHOULD prioritize information for smaller-SvcPriority records.

As discussed in [Section 3](#), clients MUST be able to fetch additional information that is required to use a SVCB record, if it is not included in the initial response. As a performance optimization, if some of the SVCB records in the response can be used without requiring additional DNS queries, the client MAY prefer those records, regardless of their priorities.

## 6. Initial SvcParamKeys

A few initial SvcParamKeys are defined here. These keys are useful for HTTPS, and most are applicable to other protocols as well.

### 6.1. "alpn" and "no-default-alpn"

The "alpn" and "no-default-alpn" SvcParamKeys together indicate the set of Application Layer Protocol Negotiation (ALPN) protocol identifiers [[ALPN](#)] and associated transport protocols supported by this service endpoint.

As with Alt-Svc [[AltSvc](#)], the ALPN protocol identifier is used to identify the application protocol and associated suite of protocols supported by the endpoint (the "protocol suite"). Clients filter the set of ALPN identifiers to match the protocol suites they support, and this informs the underlying transport protocol used (such as QUIC-over-UDP or TLS-over-TCP).

ALPNs are identified by their registered "Identification Sequence" ("alpn-id"), which is a sequence of 1-255 octets.

alpn-id = 1\*255OCTET

"alpn" is a multi-valued SvcParamKey. To construct its value-list, apply the value-list decoding algorithm (Appendix A.1) to the SvcParamValue. Each decoded value in the "alpn" value-list SHALL be an "alpn-id". The value-list MUST NOT be empty.



The wire format value for "alpn" consists of at least one "alpn-id" prefixed by its length as a single octet, and these length-value pairs are concatenated to form the SvcParamValue. These pairs MUST exactly fill the SvcParamValue; otherwise, the SvcParamValue is malformed.

For "no-default-alpn", the presentation and wire format values MUST be empty.

Each scheme that uses this SvcParamKey defines a "default set" of supported ALPNs, which SHOULD NOT be empty. To determine the set of protocol suites supported by an endpoint (the "SVCB ALPN set"), the client adds the default set to the list of "alpn-id"s unless the "no-default-alpn" SvcParamKey is present. The presence of an ALPN protocol in the SVCB ALPN set indicates that this service endpoint, described by TargetName and the other parameters (e.g. "port") offers service with the protocol suite associated with this ALPN protocol.

ALPN protocol names that do not uniquely identify a protocol suite (e.g. an Identification Sequence that can be used with both TLS and DTLS) are not compatible with this SvcParamKey and MUST NOT be included in the SVCB ALPN set.

To establish a connection to the endpoint, clients MUST

1. Let SVCB-ALPN-Intersection be the set of protocols in the SVCB ALPN set that the client supports.
2. Let Intersection-Transports be the set of transports (e.g. TLS, DTLS, QUIC) implied by the protocols in SVCB-ALPN-Intersection.
3. For each transport in Intersection-Transports, construct a ProtocolNameList containing the Identification Sequences of all the client's supported ALPN protocols for that transport, without regard to the SVCB ALPN set.

For example, if the SVCB ALPN set is ["http/1.1", "h3"], and the client supports HTTP/1.1, HTTP/2, and HTTP/3, the client could attempt to connect using TLS over TCP with a ProtocolNameList of ["http/1.1", "h2"], and could also attempt a connection using QUIC, with a ProtocolNameList of ["h3"].

Once the client has constructed a ClientHello, protocol negotiation in that handshake proceeds as specified in [\[ALPN\]](#), without regard to the SVCB ALPN set.

With this procedure in place, an attacker who can modify DNS and network traffic can prevent a successful transport connection, but



cannot otherwise interfere with ALPN protocol selection. This procedure also ensures that each ProtocolNameList includes at least one protocol from the SVCB ALPN set.

Clients SHOULD NOT attempt connection to a service endpoint whose SVCB ALPN set does not contain any supported protocols. To ensure consistency of behavior, clients MAY reject the entire SVCB RRSset and fall back to basic connection establishment if all of the RRs indicate "no-default-alpn", even if connection could have succeeded using a non-default alpn.

For compatibility with clients that require default transports, zone operators SHOULD ensure that at least one RR in each RRSset supports the default transports.

## 6.2. "port"

The "port" SvcParamKey defines the TCP or UDP port that should be used to reach this alternative endpoint. If this key is not present, clients SHALL use the authority endpoint's port number.

The presentation "value" of the SvcParamValue is a single decimal integer between 0 and 65535 in ASCII. Any other "value" (e.g. an empty value) is a syntax error. To enable simpler parsing, this SvcParam MUST NOT contain escape sequences.

The wire format of the SvcParamValue is the corresponding 2 octet numeric value in network byte order.

If a port-restricting firewall is in place between some client and the service endpoint, changing the port number might cause that client to lose access to the service, so operators should exercise caution when using this SvcParamKey to specify a non-default port.

## 6.3. "echconfig"

The SvcParamKey to enable Encrypted ClientHello (ECH) is "echconfig". Its value is defined in [Section 9](#). It is applicable to most TLS-based protocols.

When publishing a record containing an "echconfig" parameter, the publisher MUST ensure that all IP addresses of TargetName correspond to servers that have access to the corresponding private key or are authoritative for the public name. (See Section 7.2.2 of [\[ECH\]](#) for more details about the public name.) This yields an anonymity set of cardinality equal to the number of ECH-enabled server domains supported by a given client-facing server. Thus, even with an encrypted ClientHello, an attacker who can enumerate the set of ECH-



enabled domains supported by a client-facing server can guess the correct SNI with probability at least  $1/K$ , where  $K$  is the size of this ECH-enabled server anonymity set. This probability may be increased via traffic analysis or other mechanisms.

#### 6.4. "ipv4hint" and "ipv6hint"

The "ipv4hint" and "ipv6hint" keys convey IP addresses that clients MAY use to reach the service. If A and AAAA records for TargetName are locally available, the client SHOULD ignore these hints. Otherwise, clients SHOULD perform A and/or AAAA queries for TargetName as in [Section 3](#), and clients SHOULD use the IP address in those responses for future connections. Clients MAY opt to terminate any connections using the addresses in hints and instead switch to the addresses in response to the TargetName query. Failure to use A and/or AAAA response addresses could negatively impact load balancing or other geo-aware features and thereby degrade client performance.

To construct the value-list, apply the value-list decoding algorithm (Appendix A.1) to the SvcParamValue. Each decoded value in the value-list SHALL be an IP address of the appropriate family in standard textual format [[RFC5952](#)]. To enable simpler parsing, this SvcParamValue MUST NOT contain escape sequences.

The wire format for each parameter is a sequence of IP addresses in network byte order. Like an A or AAAA RSet, the list of addresses represents an unordered collection, and clients SHOULD pick addresses to use in a random order. An empty list of addresses is invalid.

When selecting between IPv4 and IPv6 addresses to use, clients may use an approach such as Happy Eyeballs [[HappyEyeballsV2](#)]. When only "ipv4hint" is present, IPv6-only clients may synthesize IPv6 addresses as specified in [[RFC7050](#)] or ignore the "ipv4hint" key and wait for AAAA resolution ([Section 3](#)). Recursive resolvers MUST NOT perform DNS64 ([[RFC6147](#)]) on parameters within a SVCB record. For best performance, server operators SHOULD include an "ipv6hint" parameter whenever they include an "ipv4hint" parameter.

These parameters are intended to minimize additional connection latency when a recursive resolver is not compliant with the requirements in [Section 4](#), and SHOULD NOT be included if most clients are using compliant recursive resolvers. When TargetName is the origin hostname or the owner name (which can be written as "."), server operators SHOULD NOT include these hints, because they are unlikely to convey any performance benefit.





## 7. ServiceMode RR compatibility and mandatory keys

In a ServiceMode RR, a SvcParamKey is considered "mandatory" if the RR will not function correctly for clients that ignore this SvcParamKey. Each SVCB protocol mapping SHOULD specify a set of keys that are "automatically mandatory", i.e. mandatory if they are present in an RR. The SvcParamKey "mandatory" is used to indicate any mandatory keys for this RR, in addition to any automatically mandatory keys that are present.

A ServiceMode RR is considered "compatible" with a client if the client recognizes all the mandatory keys, and their values indicate that successful connection establishment is possible. If the SVCB RRSets contains no compatible RRs, the client will generally act as if the RRSets is empty.

In presentation format, "mandatory" contains a list of one or more valid SvcParamKeys, either by their registered name or in the unknown-key format ([Section 2.1](#)). Keys MAY appear in any order, but MUST NOT appear more than once. Any listed keys MUST also appear in the SvcParams.

To construct the value-list, apply the value-list decoding algorithm (Appendix A.1) to the SvcParamValue. To enable simpler parsing, this SvcParamValue MUST NOT contain escape sequences.

For example, the following is a valid list of SvcParams:

```
echconfig=... key65333=ex1 key65444=ex2 mandatory=key65444,echconfig
```

In wire format, the keys are represented by their numeric values in network byte order, concatenated in ascending order.

This SvcParamKey is always automatically mandatory, and MUST NOT appear in its own value-list. Other automatically mandatory keys SHOULD NOT appear in the list either. (Including them wastes space and otherwise has no effect.)

## 8. Using SVCB with HTTPS and HTTP

Use of any protocol with SVCB requires a protocol-specific mapping specification. This section specifies the mapping for HTTPS and HTTP.

To enable special handling for the HTTPS and HTTP use-cases, the HTTPS RR type is defined as a SVCB-compatible RR type, specific to the https and http schemes. Clients MUST NOT perform SVCB queries or accept SVCB responses for "https" or "http" schemes.



The HTTPS RR wire format and presentation format are identical to SVCB, and both share the SvcParamKey registry. SVCB semantics apply equally to HTTPS RRs unless specified otherwise. The presentation format of the record is:

```
Name TTL IN HTTPS SvcPriority TargetName SvcParams
```

As with SVCB, the record is defined specifically within the Internet ("IN") Class [[RFC1035](#)].

All the SvcParamKeys defined in [Section 6](#) are permitted for use in HTTPS RRs. The default set of ALPN IDs is the single value "http/1.1". The "automatically mandatory" keys ([Section 7](#)) are "port", "alpn", and "no-default-alpn".

The presence of an HTTPS RR for an origin also indicates that all HTTP resources are available over HTTPS, as discussed in [Section 8.5](#). This allows HTTPS RRs to apply to pre-existing "http" scheme URLs, while ensuring that the client uses a secure and authenticated HTTPS connection.

The HTTPS RR parallels the concepts introduced in the HTTP Alternative Services proposed standard [[AltSvc](#)]. Clients and servers that implement HTTPS RRs are not required to implement Alt-Svc.

### **[8.1](#). Query names for HTTPS RRs**

The HTTPS RR uses Port Prefix Naming ([Section 2.3](#)), with one modification: if the scheme is "https" and the port is 443, then the client's original QNAME is equal to the service name (i.e. the origin's hostname), without any prefix labels.

By removing the Attrleaf labels [[Attrleaf](#)] used in SVCB, this construction enables offline DNSSEC signing of wildcard domains, which are commonly used with HTTPS. Reusing the service name also allows the targets of existing CNAME chains (e.g. CDN hosts) to start returning HTTPS RR responses without requiring origin domains to configure and maintain an additional delegation.

Following of HTTPS AliasMode RRs and CNAME aliases is unchanged from SVCB.

Clients always convert "http" URLs to "https" before performing an HTTPS RR query using the process described in [Section 8.5](#), so domain owners MUST NOT publish HTTPS RRs with a prefix of "\_http".



Note that none of these forms alter the HTTPS origin or authority. For example, clients **MUST** continue to validate TLS certificate hostnames based on the origin.

## **8.2. Relationship to Alt-Svc**

Publishing a ServiceMode HTTPS RR in DNS is intended to be similar to transmitting an Alt-Svc field value over HTTPS, and receiving an HTTPS RR is intended to be similar to receiving that field value over HTTPS. However, there are some differences in the intended client and server behavior.

### **8.2.1. ALPN usage**

Unlike Alt-Svc Field Values, HTTPS RRs can contain multiple ALPN IDs, and clients are encouraged to offer additional ALPNs that they support (subject to security constraints).

### **8.2.2. Untrusted channel**

SVCB does not require or provide any assurance of authenticity. (DNSSEC signing and verification, which would provide such assurance, are **OPTIONAL**.) The DNS resolution process is treated as an untrusted channel that learns only the QNAME, and is prevented from mounting any attack beyond denial of service.

Alt-Svc parameters that cannot be safely received in this model **MUST NOT** have a corresponding defined SvcParamKey. For example, there is no SvcParamKey corresponding to the Alt-Svc "persist" parameter, because this parameter is not safe to accept over an untrusted channel.

### **8.2.3. Cache lifetime**

There is no SvcParamKey corresponding to the Alt-Svc "ma" (max age) parameter. Instead, server operators encode the expiration time in the DNS TTL.

The appropriate TTL value might be different from the "ma" value used for Alt-Svc, depending on the desired efficiency and agility. Some DNS caches incorrectly extend the lifetime of DNS records beyond the stated TTL, so server operators cannot rely on HTTPS RRs expiring on time. Shortening the TTL to compensate for incorrect caching is **NOT RECOMMENDED**, as this practice impairs the performance of correctly functioning caches and does not guarantee faster expiration from incorrect caches. Instead, server operators **SHOULD** maintain compatibility with expired records until they observe that nearly all connections have migrated to the new configuration.



#### **8.2.4. Granularity**

Sending Alt-Svc over HTTP allows the server to tailor the Alt-Svc Field Value specifically to the client. When using an HTTPS RR, groups of clients will necessarily receive the same SvcParams. Therefore, HTTPS RRs are not suitable for uses that require single-client granularity.

#### **8.3. Interaction with Alt-Svc**

Clients that do not implement support for Encrypted ClientHello MAY skip the HTTPS RR query if a usable Alt-Svc value is available in the local cache. If Alt-Svc connection fails, these clients SHOULD fall back to the HTTPS RR client connection procedure ([Section 3](#)).

For clients that implement support for ECH, the interaction between HTTPS RRs and Alt-Svc is described in [Section 9.1](#).

This specification does not alter the DNS queries performed when connecting to an Alt-Svc hostname (typically A and/or AAAA only).

#### **8.4. Requiring Server Name Indication**

Clients MUST NOT use an HTTPS RR response unless the client supports TLS Server Name Indication (SNI) and indicate the origin name when negotiating TLS. This supports the conservation of IP addresses.

Note that the TLS SNI (and also the HTTP "Host" or ":authority") will indicate the origin, not the TargetName.

#### **8.5. HTTP Strict Transport Security**

By publishing a usable HTTPS RR, the server operator indicates that all useful HTTP resources on that origin are reachable over HTTPS, similar to HTTP Strict Transport Security [[HSTS](#)].

Prior to making an "http" scheme request, the client SHOULD perform a lookup to determine if any HTTPS RRs exist for that origin. To do so, the client SHOULD construct a corresponding "https" URL as follows:

1. Replace the "http" scheme with "https".
2. If the "http" URL explicitly specifies port 80, specify port 443.
3. Do not alter any other aspect of the URL.

This construction is equivalent to Section 8.3 of [[HSTS](#)], point 5.





If an HTTPS RR query for this "https" URL returns any AliasMode HTTPS RRs, or any compatible ServiceMode HTTPS RRs (see [Section 7](#)), the client SHOULD act as if it has received an HTTP "307 Temporary Redirect" redirect to this "https" URL. (Receipt of an incompatible ServiceMode RR does not trigger the redirect behavior.) Because HTTPS RRs are received over an often insecure channel (DNS), clients MUST NOT place any more trust in this signal than if they had received a 307 redirect over cleartext HTTP.

When making an "https" scheme request to an origin with an HTTPS RR, either directly or via the above redirect, the client SHOULD terminate the connection if there are any errors with the underlying secure transport, such as errors in certificate validation. This aligns with [Section 8.4](#) and Section 12.1 of [\[HSTS\]](#).

### **8.6. HTTP-based protocols**

All protocols employing "http://" or "https://" URLs SHOULD respect HTTPS RRs. For example, clients that support HTTPS RRs and implement the altered WebSocket [\[WebSocket\]](#) opening handshake from the W3C Fetch specification [\[FETCH\]](#) SHOULD use HTTPS RRs for the "requestURL".

An HTTP-based protocol MAY define its own SVCB mapping. Such mappings MAY be defined to take precedence over HTTPS RRs.

## **9. SVCB/HTTPS RR parameter for ECH configuration**

The SVCB "echconfig" parameter is defined for conveying the ECH configuration of an alternative endpoint. In wire format, the value of the parameter is an ECHConfigs vector [\[ECH\]](#), including the redundant length prefix. In presentation format, the value is a single ECHConfigs encoded in Base64 [\[base64\]](#). Base64 is used here to simplify integration with TLS server software. To enable simpler parsing, this SvcParam MUST NOT contain escape sequences.

When ECH is in use, the TLS ClientHello is divided into an unencrypted "outer" and an encrypted "inner" ClientHello. The outer ClientHello is an implementation detail of ECH, and its contents are controlled by the ECHConfig in accordance with [\[ECH\]](#). The inner ClientHello is used for establishing a connection to the service, so its contents may be influenced by other SVCB parameters. For example, the requirements on the ProtocolNameList in [Section 6.1](#) apply only to the inner ClientHello. Similarly, it is the inner ClientHello whose Server Name Indication identifies the desired service.



### **9.1. Client behavior**

The general client behavior specified in [Section 3](#) permits clients to retry connection with a less preferred alternative if the preferred option fails, including falling back to a direct connection if all SVCB options fail. This behavior is not suitable for ECH, because fallback would negate the privacy benefits of ECH. Accordingly, ECH-capable clients SHALL implement the following behavior for connection establishment:

1. Perform connection establishment using HTTPS RRs as described in [Section 3](#). After step 4, if there were compatible HTTPS RRs, they all had an "echconfig" key, and attempts to connect to them all failed, terminate connection establishment.
2. If the client implements Alt-Svc, try to connect using any entries from the Alt-Svc cache.
3. Continue connection establishment as in [Section 3](#) if necessary.

As a latency optimization, clients MAY prefetch DNS records for later steps before they are needed.

### **9.2. Deployment considerations**

An HTTPS RRSet containing some RRs with "echconfig" and some without is vulnerable to a downgrade attack. This configuration is NOT RECOMMENDED. Zone owners who do use such a mixed configuration SHOULD mark the RRs with "echconfig" as more preferred (i.e. smaller SvcPriority) than those without, in order to maximize the likelihood that ECH will be used in the absence of an active adversary.

## **10. Zone Structures**

### **10.1. Structuring zones for flexibility**

Each ServiceForm RRSet can only serve a single scheme. The scheme is indicated by the owner name and the RR type. For the generic SVCB RR type, this means that each owner name can only be used for a single scheme. The underscore prefixing requirement ([Section 2.3](#)) ensures that this is true for the initial query, but it is the responsibility of zone owners to choose names that satisfy this constraint when using aliases, including CNAME and AliasMode records.

When using the generic SVCB RR type with aliasing, zone owners SHOULD choose alias target names that indicate the scheme in use (e.g. "foosvc.example.net" for "foo://" schemes). This will help to avoid confusion when another scheme needs to be added to the configuration.



## [10.2.](#) Structuring zones for performance

To avoid a delay for clients using a nonconforming recursive resolver, domain owners SHOULD minimize the use of AliasMode records, and choose TargetName to be a domain for which the client will have already issued address queries (see [Section 3](#)). For foo://foo.example.com:8080, this might look like:

```
$ORIGIN example.com. ; Origin
foo                3600 IN CNAME foosvc.example.net.
_8080._foo.foo     3600 IN CNAME foosvc.example.net.

$ORIGIN example.net. ; Service provider zone
foosvc             3600 IN SVCB 1 . key65333=...
foosvc             300  IN AAAA 2001:db8::1
```

Domain owners SHOULD avoid using a TargetName that is below a DNAME, as this is likely unnecessary and makes responses slower and larger.

## [10.3.](#) Examples

### [10.3.1.](#) Protocol enhancements

Consider a simple zone of the form:

```
$ORIGIN simple.example. ; Simple example zone
@ 300 IN A      192.0.2.1
      AAAA 2001:db8::1
```

The domain owner could add this record:

```
simple.example. 7200 IN HTTPS 1 . alpn=h3
```

to indicate that simple.example uses HTTPS, and supports QUIC in addition to HTTPS over TCP (an implicit default). The record could also include other information (e.g. non-standard port, ECH configuration).

### [10.3.2.](#) Apex aliasing

Consider a zone that is using CNAME aliasing:

```
$ORIGIN aliased.example. ; A zone that is using a hosting service
; Subdomain aliased to a high-performance server pool
www                7200 IN CNAME pool.svc.example.
; Apex domain on fixed IPs because CNAME is not allowed at the apex
@                  300  IN A      192.0.2.1
                  IN AAAA 2001:db8::1
```



With HTTPS RRs, the owner of `aliased.example` could alias the apex by adding one additional record:

```
@           7200 IN HTTPS 0 pool.svc.example.
```

With this record in place, HTTPS-RR-aware clients will use the same server pool for `aliased.example` and `www.aliased.example`. (They will also upgrade to HTTPS on `aliased.example`.) Non-HTTPS-RR-aware clients will just ignore the new record.

Similar to CNAME, HTTPS RRs have no impact on the origin name. When connecting, clients will continue to treat the authoritative origins as "`https://www.aliased.example`" and "`https://aliased.example`", respectively, and will validate TLS server certificates accordingly.

#### **10.3.3. Parameter binding**

Suppose that `svc.example`'s default server pool supports HTTP/2, and it has deployed HTTP/3 on a new server pool with a different configuration. This can be expressed in the following form:

```
$ORIGIN svc.example. ; A hosting provider.
pool  7200 IN HTTPS 1 h3pool alpn=h2,h3 echconfig="123..."
                HTTPS 2 .      alpn=h2 echconfig="abc..."
pool   300 IN A        192.0.2.2
                AAAA    2001:db8::2
h3pool 300 IN A        192.0.2.3
                AAAA    2001:db8::3
```

This configuration is entirely compatible with the "Apex aliasing" example, whether the client supports HTTPS RRs or not. If the client does support HTTPS RRs, all connections will be upgraded to HTTPS, and clients will use HTTP/3 if they can. Parameters are "bound" to each server pool, so each server pool can have its own protocol, ECH configuration, etc.

#### **10.3.4. Multi-CDN**

The HTTPS RR is intended to support HTTPS services operated by multiple independent entities, such as different Content Delivery Networks (CDNs) or different hosting providers. This includes the case where a service is migrated from one operator to another, as well as the case where the service is multiplexed between multiple operators for performance, redundancy, etc.

This example shows such a configuration, with `www.customer.example` having different DNS responses to different queries, either over time or due to logic within the authoritative DNS server:





```
; This zone contains/returns different CNAME records
; at different points-in-time. The RRset for "www" can
; only ever contain a single CNAME.

; Sometimes the zone has:
$ORIGIN customer.example. ; A Multi-CDN customer domain
www 900 IN CNAME cdn1.svc1.example.

; and other times it contains:
$ORIGIN customer.example.
www 900 IN CNAME customer.svc2.example.

; and yet other times it contains:
$ORIGIN customer.example.
www 900 IN CNAME cdn3.svc3.example.

; With the following remaining constant and always included:
$ORIGIN customer.example. ; A Multi-CDN customer domain
; The apex is also aliased to www to match its configuration
@      7200 IN HTTPS 0 www
; Non-HTTPS-aware clients use non-CDN IPs
      A      203.0.113.82
      AAAA 2001:db8:203::2

; Resolutions following the cdn1.svc1.example
; path use these records.
; This CDN uses a different alternative service for HTTP/3.
$ORIGIN svc1.example. ; domain for CDN 1
cdn1      1800 IN HTTPS 1 h3pool alpn=h3 echconfig="123..."
          HTTPS 2 . alpn=h2 echconfig="123..."
          A      192.0.2.2
          AAAA 2001:db8:192::4
h3pool 300 IN A 192.0.2.3
          AAAA 2001:db8:192:7::3

; Resolutions following the customer.svc2.example
; path use these records.
; Note that this CDN only supports HTTP/2.
$ORIGIN svc2.example. ; domain operated by CDN 2
customer 300 IN HTTPS 1 . alpn=h2 echconfig="xyz..."
        60 IN A      198.51.100.2
          A      198.51.100.3
          A      198.51.100.4
          AAAA 2001:db8:198::7
          AAAA 2001:db8:198::12

; Resolutions following the customer.svc2.example
; path use these records.
```



```
; Note that this CDN has no HTTPS records
; and thus no ECH support.
$ORIGIN svc3.example. ; domain operated by CDN 3
cdn3      60 IN A      203.0.113.8
          AAAA 2001:db8:113::8
```

Note that in the above example, the different CDNs have different echconfig and different capabilities, but clients will use HTTPS RRs as a bound-together unit.

Domain owners should be cautious when using a multi-CDN configuration, as it introduces a number of complexities highlighted by this example:

- o If CDN 1 supports ECH, and CDN 2 does not, the client is vulnerable to ECH downgrade by a network adversary who forces clients to get CDN 2 records.
- o Aliasing the apex to its subdomain simplifies the zone file but likely increases resolution latency, especially when using a non-HTTPS-aware recursive resolver. An alternative would be to alias the zone apex directly to a name managed by a CDN.
- o The A, AAAA, HTTPS resolutions are independent lookups so clients may observe and follow different CNAMEs to different CDNs. Clients may thus find a SvcDomainName pointing to a name other than the one which returned along with the A and AAAA lookups and will need to do an additional resolution for them. Including ipv6hint and ipv4hint will reduce the performance impact of this case.
- o If not all CDNs publish HTTPS records, clients will sometimes receive NODATA for HTTPS queries (as with cdn3.svc3.example above), and thus no echconfig, but could receive A/AAAA records from a different CDN which does support ECH. Clients will be unable to use ECH in this case.

#### **10.3.5. Non-HTTPS uses**

For services other than HTTPS, the SVCB RR and an Attrleaf label [[Attrleaf](#)] will be used. For example, to reach an example resource of "baz://api.example.com:8765", the following SVCB record would be used to alias it to "svc4-baz.example.net." which in-turn could return AAAA/A records and/or SVCB records in ServiceMode:

```
_8765._baz.api.example.com. 7200 IN SVCB 0 svc4-baz.example.net.
```



HTTPS RRs use similar Attrleaf labels if the origin contains a non-default port.

## **11. Interaction with other standards**

This standard is intended to reduce connection latency and improve user privacy. Server operators implementing this standard SHOULD also implement TLS 1.3 [[RFC8446](#)] and OCSP Stapling [[RFC6066](#)], both of which confer substantial performance and privacy benefits when used in combination with SVCB records.

To realize the greatest privacy benefits, this proposal is intended for use over a privacy-preserving DNS transport (like DNS over TLS [[DoT](#)] or DNS over HTTPS [[DoH](#)]). However, performance improvements, and some modest privacy improvements, are possible without the use of those standards.

Any specification for use of SVCB with a protocol MUST have an entry for its scheme under the SVCB RR type in the IANA DNS Underscore Global Scoped Entry Registry [[Attrleaf](#)]. The scheme SHOULD have an entry in the IANA URI Schemes Registry [[RFC7595](#)]. The scheme SHOULD have a defined specification for use with SVCB.

## **12. Security Considerations**

SVCB/HTTPS RRs are intended for distribution over untrusted channels, and clients are REQUIRED to verify that the alternative endpoint is authoritative for the service (similar to Section 2.1 of [[AltSvc](#)]). Therefore, DNSSEC signing and validation are OPTIONAL for publishing and using SVCB and HTTPS RRs.

Clients MUST ensure that their DNS cache is partitioned for each local network, or flushed on network changes, to prevent a local adversary in one network from implanting a forged DNS record that allows them to track users or hinder their connections after they leave that network.

An attacker who can prevent SVCB resolution can deny clients any associated security benefits. A hostile recursive resolver can always deny service to SVCB queries, but network intermediaries can often prevent resolution as well, even when the client and recursive resolver validate DNSSEC and use a secure transport. These downgrade attacks can prevent the HTTPS upgrade provided by the HTTPS RR ([Section 8.5](#)), and disable the encryption enabled by the echconfig SvcParamKey ([Section 9](#)). To prevent downgrades, [Section 3.1](#) recommends that clients abandon the connection attempt when such an attack is detected.



A hostile DNS intermediary might forge AliasForm "." records ([Section 2.5.1](#)) as a way to block clients from accessing particular services. Such an adversary could already block entire domains by forging erroneous responses, but this mechanism allows them to target particular protocols or ports within a domain. Clients that might be subject to such attacks SHOULD ignore AliasForm "." records.

### **[13.](#) Privacy Considerations**

Standard address queries reveal the user's intent to access a particular domain. This information is visible to the recursive resolver, and to many other parties when plaintext DNS transport is used. SVCB queries, like queries for SRV records and other specific RR types, additionally reveal the user's intent to use a particular protocol. This is not normally sensitive information, but it should be considered when adding SVCB support in a new context.

### **[14.](#) IANA Considerations**

#### **[14.1.](#) SVCB RRTYPE**

This document defines a new DNS RR type, SVCB, whose value 64 has been allocated by IANA from the "Resource Record (RR) TYPES" subregistry of the "Domain Name System (DNS) Parameters" registry:

Type: SVCB

Value: 64

Meaning: General Purpose Service Endpoints

Reference: This document

#### **[14.2.](#) HTTPS RRTYPE**

This document defines a new DNS RR type, HTTPS, whose value 65 has been allocated by IANA from the "Resource Record (RR) TYPES" subregistry of the "Domain Name System (DNS) Parameters" registry:

Type: HTTPS

Value: 65

Meaning: HTTPS Specific Service Endpoints

Reference: This document





### **14.3. New registry for Service Parameters**

The "Service Binding (SVCB) Parameter Registry" defines the namespace for parameters, including string representations and numeric SvcParamKey values. This registry is shared with other SVCB-compatible RR types, such as the HTTPS RR.

ACTION: create and include a reference to this registry.

#### **14.3.1. Procedure**

A registration MUST include the following fields:

- o Number: SvcParamKey wire format numeric identifier (range 0-65535)
- o Name: SvcParamKey presentation name
- o Meaning: a short description
- o Pointer to specification text

SvcParamKey entries to be added to this namespace have different policies ([\[RFC8126\]](#), [Section 4](#)) based on their range:

Number	IANA Policy
0-255	Standards Action
256-32767	Expert Review
32768-65280	First Come First Served
65280-65534	Private Use
65535	Standards Action

Apart from the initial contents, the SvcParamKey name MUST NOT start with "key".

#### **14.3.2. Initial contents**

The "Service Binding (SVCB) Parameter Registry" shall initially be populated with the registrations below:



Number	Name	Meaning	Reference
0	mandatory	Mandatory keys in this RR	(This document)
1	alpn	Additional supported protocols	(This document)
2	no-default-alpn	No support for default protocol	(This document)
3	port	Port for alternative endpoint	(This document)
4	ipv4hint	IPv4 address hints	(This document)
5	echconfig	Encrypted ClientHello info	(This document)
6	ipv6hint	IPv6 address hints	(This document)
65280-65534	keyNNNNN	Private Use	(This document)
65535	key65535	Reserved ("Invalid key")	(This document)

#### 14.4. Registry updates

Per [\[RFC6895\]](#), please add the following entries to the data type range of the Resource Record (RR) TYPEs registry:

TYPE	Meaning	Reference
SVCB	Service Location and Parameter Binding	(This document)
HTTPS	HTTPS Service Location and Parameter Binding	(This document)

Per [\[Attrleaf\]](#), please add the following entry to the DNS Underscore Global Scoped Entry Registry:



RR TYPE	_NODE NAME	Meaning	Reference
HTTPS	_https	HTTPS SVCB info	(This document)

## 15. Acknowledgments and Related Proposals

There have been a wide range of proposed solutions over the years to the "CNAME at the Zone Apex" challenge proposed. These include [[I-D.bellis-dnsop-http-record](#)], [[I-D.ietf-dnsop-aname](#)], and others.

Thank you to Ian Swett, Ralf Weber, Jon Reed, Martin Thomson, Lucas Pardue, Ilari Liusvaara, Tim Wicinski, Tommy Pauly, Chris Wood, David Benjamin, Mark Andrews, Emily Stark, Eric Orth, Kyle Rose, Craig Taylor, Dan McArdle, Brian Dickson, and others for their feedback and suggestions on this draft.

## 16. References

### 16.1. Normative References

- [ALPN] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", [RFC 7301](#), DOI 10.17487/RFC7301, July 2014, <<https://www.rfc-editor.org/info/rfc7301>>.
- [Attrleaf] Crocker, D., "Scoped Interpretation of DNS Resource Records through "Underscored" Naming of Attribute Leaves", [BCP 222](#), [RFC 8552](#), DOI 10.17487/RFC8552, March 2019, <<https://www.rfc-editor.org/info/rfc8552>>.
- [base64] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", [RFC 4648](#), DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [DNAME] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", [RFC 6672](#), DOI 10.17487/RFC6672, June 2012, <<https://www.rfc-editor.org/info/rfc6672>>.
- [DoH] 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>>.



- [DoT] 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>>.
- [ECH] Rescorla, E., Oku, K., Sullivan, N., and C. Wood, "TLS Encrypted Client Hello", [draft-ietf-tls-esni-08](#) (work in progress), October 2020.
- [HappyEyeballsV2] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", [RFC 8305](#), DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [HSTS] Hodges, J., Jackson, C., and A. Barth, "HTTP Strict Transport Security (HSTS)", [RFC 6797](#), DOI 10.17487/RFC6797, November 2012, <<https://www.rfc-editor.org/info/rfc6797>>.
- [HTTP3] Bishop, M., "Hypertext Transfer Protocol Version 3 (HTTP/3)", [draft-ietf-quic-http-32](#) (work in progress), October 2020.
- [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>>.
- [RFC1928] Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5", [RFC 1928](#), DOI 10.17487/RFC1928, March 1996, <<https://www.rfc-editor.org/info/rfc1928>>.
- [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>>.
- [RFC3225] Conrad, D., "Indicating Resolver Support of DNSSEC", [RFC 3225](#), DOI 10.17487/RFC3225, December 2001, <<https://www.rfc-editor.org/info/rfc3225>>.
- [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>>.





- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", [RFC 5952](#), DOI 10.17487/RFC5952, August 2010, <<https://www.rfc-editor.org/info/rfc5952>>.
- [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>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", [RFC 7050](#), DOI 10.17487/RFC7050, November 2013, <<https://www.rfc-editor.org/info/rfc7050>>.
- [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>>.
- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", [BCP 35](#), [RFC 7595](#), DOI 10.17487/RFC7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", [RFC 8446](#), DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.



[WebSocket]

Fette, I. and A. Melnikov, "The WebSocket Protocol", [RFC 6455](#), DOI 10.17487/RFC6455, December 2011, <<https://www.rfc-editor.org/info/rfc6455>>.

## 16.2. Informative 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>>.
- [DNSTerm] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", [BCP 219](#), [RFC 8499](#), DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [FETCH] "Fetch Living Standard", May 2020, <<https://fetch.spec.whatwg.org/>>.
- [I-D.bellis-dnsop-http-record] Bellis, R., "A DNS Resource Record for HTTP", [draft-bellis-dnsop-http-record-00](#) (work in progress), November 2018.
- [I-D.ietf-dnsop-aname] Finch, T., Hunt, E., Dijk, P., Eden, A., and W. Mekking, "Address-specific DNS aliases (ANAME)", [draft-ietf-dnsop-aname-04](#) (work in progress), July 2019.
- [RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/info/rfc6454>>.
- [RFC6895] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", [BCP 42](#), [RFC 6895](#), DOI 10.17487/RFC6895, April 2013, <<https://www.rfc-editor.org/info/rfc6895>>.
- [SRV] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), DOI 10.17487/RFC2782, February 2000, <<https://www.rfc-editor.org/info/rfc2782>>.
- [URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.



### 16.3. URIs

[1] <https://github.com/MikeBishop/dns-alt-svc>

### Appendix A. Decoding text in zone files

DNS zone files are capable of representing arbitrary octet sequences in basic ASCII text, using various delimiters and encodings. The algorithm for decoding these character-strings is defined in [Section 5.1 of \[RFC1035\]](#). Here we summarize the allowed input to that algorithm, using ABNF:

```
; non-special is VCHAR minus DQUOTE, ";", "(", ")", and "\".
non-special = %x21 / %x23-27 / %x2A-3A / %x3C-5B / %x5D-7E
; non-digit is VCHAR minus DIGIT
non-digit   = %x21-2F / %x3A-7E
; dec-octet is a number 0-255 as a three-digit decimal number.
dec-octet   = ( "0" / "1" ) 2DIGIT /
               "2" ( ( %x30-34 DIGIT ) / ( "5" %x30-35 ) )
escaped     = "\" ( non-digit / dec-octet )
contiguous  = 1*( non-special / escaped )
quoted      = DQUOTE *( contiguous / ( ["\""] WSP ) ) DQUOTE
char-string = contiguous / quoted
```

The decoding algorithm allows "char-string" to represent any "\*OCTET". In this document, this algorithm is referred to as "character-string decoding". The algorithm is the same as used by "<character-string>" in [RFC 1035](#), although the output length in this document is not limited to 255 octets.

#### A.1. Decoding a value-list

In order to represent lists of values in zone files, this specification uses an extended version of character-string decoding that adds the use of "," as a delimiter after double-quote processing. When "," is not escaped (by a preceding "\" or as the escape sequence "\044"), it separates values in the output, which is a list of 1\*OCTET. (For simplicity, empty values are not allowed.) We refer to this modified procedure as "value-list decoding".

```
value-list = char-string
list-value = 1*OCTET
```

For example, consider these "char-string" SvcParamValues:

```
"part1,part2\,part3"
part1,part2\044part3
```



Character-string decoding either of these inputs would produce a single `"*OCTET"` output:

```
part1,part2,part3
```

Value-list decoding either of these inputs would instead convert it to a list of two `"list-value"`s:

```
part1  
part2,part3
```

## **[Appendix B](#). Comparison with alternatives**

The SVCB and HTTPS RR types closely resemble, and are inspired by, some existing record types and proposals. A complaint with all of the alternatives is that web clients have seemed unenthusiastic about implementing them. The hope here is that by providing an extensible solution that solves multiple problems we will overcome the inertia and have a path to achieve client implementation.

### **[B.1](#). Differences from the SRV RR type**

An SRV record [[SRV](#)] can perform a similar function to the SVCB 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 SVCB.
- o SRV records cannot instruct the client to switch or upgrade protocols, whereas SVCB can signal such an upgrade (e.g. to HTTP/2).
- o SRV records are not extensible, whereas SVCB and HTTPS RRs can be extended with new parameters.
- o SVCB records use 16 bit for SvcPriority for consistency with SRV and other RR types that also use 16 bit priorities.

### **[B.2](#). Differences from the proposed HTTP record**

Unlike [[I-D.bellis-dnsop-http-record](#)], this approach is extensible to cover Alt-Svc and Encrypted ClientHello use-cases. Like that proposal, this addresses the zone apex CNAME challenge.

Like that proposal, it remains necessary to continue to include address records at the zone apex for legacy clients.





### **B.3. Differences from the proposed ANAME record**

Unlike [[I-D.ietf-dnsop-aname](#)], this approach is extensible to cover Alt-Svc and ECH use-cases. This approach also does not require any changes or special handling on either authoritative or primary servers, beyond optionally returning in-bailiwick additional records.

Like that proposal, this addresses the zone apex CNAME challenge for clients that implement this.

However, with this SVCB proposal, it remains necessary to continue to include address records at the zone apex for legacy clients. If deployment of this standard is successful, the number of legacy clients will fall over time. As the number of legacy clients declines, the operational effort required to serve these users without the benefit of SVCB indirection should fall. Server operators can easily observe how much traffic reaches this legacy endpoint, and may remove the apex's address records if the observed legacy traffic has fallen to negligible levels.

### **B.4. Comparison with separate RR types for AliasMode and ServiceMode**

Abstractly, functions of AliasMode and ServiceMode are independent, so it might be tempting to specify them as separate RR types. However, this would result in a serious performance impairment, because clients cannot rely on their recursive resolver to follow SVCB aliases (unlike CNAME). Thus, clients would have to issue queries for both RR types in parallel, potentially at each step of the alias chain. Recursive resolvers that implement the specification would, upon receipt of a ServiceMode query, emit both a ServiceMode and an AliasMode query to the authoritative. Thus, splitting the RR type would double, or in some cases triple, the load on clients and servers, and would not reduce implementation complexity.

## **Appendix C. Change history**

- o [draft-ietf-dnsop-svcb-https-02](#)
  - \* Added a Privacy Considerations section
  - \* Adjusted resolution fallback description
  - \* Clarified status of SvcParams in AliasMode
  - \* Improved advice on zone structuring and use with Alt-Svc
  - \* Improved examples, including a new Multi-CDN example



- \* Reorganized text on value-list parsing and SvcPriority
- \* Improved phrasing and other editorial improvements throughout
- o [draft-ietf-dnsop-svcb-https-01](#)
  - \* Added a "mandatory" SvcParamKey
  - \* Added the ability to indicate that a service does not exist
  - \* Adjusted resolution and ALPN algorithms
  - \* Major terminology revisions for "origin" and CamelCase names
  - \* Revised ABNF
  - \* Include allocated RR type numbers
  - \* Various corrections, explanations, and recommendations
- o [draft-ietf-dnsop-svcb-https-00](#)
  - \* Rename HTTPSSVC RR to HTTPS RR
  - \* Rename "an SVCB" to "a SVCB"
  - \* Removed "design considerations and open issues" section and some other "to be removed" text
- o [draft-ietf-dnsop-svcb-httpssvc-03](#)
  - \* Revised chain length limit requirements
  - \* Revised IANA registry rules for SvcParamKeys
  - \* Require HTTPS clients to implement SNI
  - \* Update terminology for Encrypted ClientHello
  - \* Clarifications: non-default ports, transport proxies, HSTS procedure, WebSocket behavior, wire format, IP hints, inner/outer ClientHello with ECH
  - \* Various textual and ABNF corrections
- o [draft-ietf-dnsop-svcb-httpssvc-02](#)
  - \* All changes to Alt-Svc have been removed



- \* Expanded and reorganized examples
  - \* Priority zero is now the definition of AliasForm
  - \* Repeated SvcParamKeys are no longer allowed
  - \* The "=" sign may be omitted in a key=value pair if the value is also empty
  - \* In the wire format, SvcParamKeys must be in sorted order
  - \* New text regarding how to handle resolution timeouts
  - \* Expanded description of recursive resolver behavior
  - \* Much more precise description of the intended ALPN behavior
  - \* Match the HSTS specification's language on HTTPS enforcement
  - \* Removed 'esniconfig=""' mechanism and simplified ESNI connection logic
- o [draft-ietf-dnsop-svcb-httpssvc-01](#)
    - \* Reduce the emphasis on conversion between HTTPSSVC and Alt-Svc
    - \* Make the "untrusted channel" concept more precise.
    - \* Make SvcFieldPriority = 0 the definition of AliasForm, instead of a requirement.
  - o [draft-ietf-dnsop-svcb-httpssvc-00](#)
    - \* Document an optimization for optimistic pre-connection. (Chris Wood)
    - \* Relax IP hint handling requirements. (Eric Rescorla)
  - o [draft-nygren-dnsop-svcb-httpssvc-00](#)
    - \* Generalize to an SVCB record, with special-case handling for Alt-Svc and HTTPS separated out to dedicated sections.
    - \* Split out a separate HTTPSSVC record for the HTTPS use-case.
    - \* Remove the explicit SvcRecordType=0/1 and instead make the AliasForm vs ServiceForm be implicit. This was based on feedback recommending against subtyping RR type.



- \* Remove one optimization.
- o [draft-nygren-httpbis-httpssvc-03](#)
  - \* Change redirect type for HSTS-style behavior from 302 to 307 to reduce ambiguities.
- o [draft-nygren-httpbis-httpssvc-02](#)
  - \* Remove the redundant length fields from the wire format.
  - \* Define a SvcDomainName of "." for SvcRecordType=1 as being the HTTPSSVC RRNAME.
  - \* Replace "hq" with "h3".
- o [draft-nygren-httpbis-httpssvc-01](#)
  - \* Fixes of record name. Replace references to "HTTPSVC" with "HTTPSSVC".
- o [draft-nygren-httpbis-httpssvc-00](#)
  - \* Initial version

#### Authors' Addresses

Ben Schwartz  
Google

Email: bemasc@google.com

Mike Bishop  
Akamai Technologies

Email: mbishop@evequefou.be

Erik Nygren  
Akamai Technologies

Email: erik+ietf@nygren.org



