
Stream: Internet Engineering Task Force (IETF)
RFC: [9461](#)
Category: Standards Track
Published: November 2023
ISSN: 2070-1721
Author: B. Schwartz
Meta Platforms, Inc.

RFC 9461

Service Binding Mapping for DNS Servers

Abstract

The SVCB DNS resource record type expresses a bound collection of endpoint metadata, for use when establishing a connection to a named service. DNS itself can be such a service, when the server is identified by a domain name. This document provides the SVCB mapping for named DNS servers, allowing them to indicate support for encrypted transport protocols.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9461>.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Definitions	3
3. Identities and Names	3
3.1. Special Case: Non-default Ports	4
4. Applicable Existing SvcParamKeys	4
4.1. "alpn"	4
4.2. "port"	4
4.3. Other Applicable SvcParamKeys	5
5. New SvcParamKey: "dohpath"	5
6. Limitations	6
7. Examples	6
8. Security Considerations	7
8.1. Adversary on the Query Path	7
8.1.1. Downgrade Attacks	7
8.1.2. Redirection Attacks	7
8.2. Adversary on the Transport Path	8
9. IANA Considerations	8
10. References	8
10.1. Normative References	8
10.2. Informative References	9
Appendix A. Mapping Summary	9
Acknowledgments	10
Author's Address	10

1. Introduction

The SVCB resource record (RR) type [SVCB] provides clients with information about how to reach alternative endpoints for a service. These endpoints may offer improved performance or privacy properties. The service is identified by a "scheme" indicating the service type, a hostname, and, optionally, other information such as a port number. A DNS server is often identified only by its IP address (e.g., in DHCP), but in some contexts it can also be identified by a hostname (e.g., "NS" records, manual resolver configuration) and sometimes also a non-default port number.

The use of the SVCB RR type requires a mapping document for each service type (Section 2.4.3 of [SVCB]), indicating how a client for that service can interpret the contents of the SVCB SvcParams. This document provides the mapping for the "dns" service type, allowing DNS servers to offer alternative endpoints and transports, including encrypted transports like DNS over TLS (DoT) [RFC7858], DNS over HTTPS (DoH) [RFC8484], and DNS over QUIC (DoQ) [RFC9250].

The SVCB mapping described in this document is intended as a general-purpose baseline. Subsequent specifications will adapt this mechanism as needed to support specific configurations (e.g., for communication between stub resolvers and recursive resolvers).

2. Conventions and Definitions

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

3. Identities and Names

SVCB record names (i.e., QNAMEs) for DNS services are formed using Port Prefix Naming (Section 2.3 of [SVCB]), with a scheme of "dns". For example, SVCB records for a DNS service identified as `dns1.example.com` would be queried at `_dns.dns1.example.com`.

In some use cases, the name used for retrieving these DNS records is different from the server identity used to authenticate the secure transport. To distinguish between these, this document uses the following terms:

Binding authority: The service name (Section 1.3 of [SVCB]) and optional port number used as input to Port Prefix Naming.

Authentication name: The name used for secure transport authentication. This **MUST** be a DNS hostname or a literal IP address. Unless otherwise specified, this is the service name from the binding authority.

3.1. Special Case: Non-default Ports

Normally, a DNS service is identified by an IP address or a domain name. When connecting to the service using unencrypted DNS over UDP or TCP, clients use the default port number for DNS (53). However, in rare cases, a DNS service might be identified by both a name and a port number. For example, the DNS URI scheme [DNSURI] optionally includes an authority, comprised of a host and a port number (with a default of 53). DNS URIs normally omit the authority or specify an IP address, but a hostname and non-default port number are allowed.

When the binding authority specifies a non-default port number, Port Prefix Naming places the port number in an additional prefix on the name. For example, if the binding authority is "dns1.example.com:9953", the client would query for SVCB records at `_9953._dns.dns1.example.com`. If two DNS services operating on different port numbers provide different behaviors, this arrangement allows them to preserve the distinction when specifying alternative endpoints.

4. Applicable Existing SvcParamKeys

4.1. "alpn"

This key indicates the set of supported protocols (Section 7.1 of [SVCB]). There is no default protocol, so the "no-default-alpn" key does not apply. If the "alpn" SvcParamKey is absent, the client **MUST** treat the SVCB record as "incompatible" (as defined in Section 8 of [SVCB]) unless some other recognized SvcParam indicates a supported protocol.

If the protocol set contains any HTTP versions (e.g., "h2", "h3"), then the record indicates support for DoH and the "dohpath" key **MUST** be present (Section 5). All keys specified for use with the HTTPS record are also permissible and apply to the resulting HTTP connection.

If the protocol set contains protocols with different default ports and no "port" key is specified, then protocols are contacted separately on their default ports. Note that in this configuration, Application-Layer Protocol Negotiation (ALPN) negotiation does not defend against cross-protocol downgrade attacks.

4.2. "port"

This key is used to indicate the target port for connection (Section 7.2 of [SVCB]). If omitted, the client **SHALL** use the default port number for each transport protocol (853 for DoT and DoQ, 443 for DoH).

This key is automatically mandatory for this binding. This means that a client that does not respect the "port" key **MUST** ignore any SVCB record that contains this key. (See Section 8 of [SVCB] for the definition of "automatically mandatory".)

Support for the "port" key can be unsafe if the client has implicit elevated access to some network service (e.g., a local service that is inaccessible to remote parties) and that service uses a TCP-based protocol other than TLS. A hostile DNS server might be able to manipulate this service by causing the client to send a specially crafted TLS Server Name Indication (SNI) or session ticket that can be misparsed as a command or exploit. To avoid such attacks, clients **SHOULD NOT** support the "port" key unless one of the following conditions applies:

- The client is being used with a DNS server that it trusts not to attempt this attack.
- The client is being used in a context where implicit elevated access cannot apply.
- The client restricts the set of allowed TCP port values to exclude any ports where a confusion attack is likely to be possible (e.g., the "bad ports" list from Section 2.9 ("Port blocking") of [FETCH]).

4.3. Other Applicable SvcParamKeys

These SvcParamKeys from [SVCB] apply to the "dns" scheme without modification:

- mandatory
- ipv4hint
- ipv6hint

Future SvcParamKeys might also be applicable.

5. New SvcParamKey: "dohpath"

"dohpath" is a single-valued SvcParamKey whose value (in both presentation format and wire format) **MUST** be a URI Template in relative form ([RFC6570], Section 1.1) encoded in UTF-8 [RFC3629]. If the "alpn" SvcParam indicates support for HTTP, "dohpath" **MUST** be present. The URI Template **MUST** contain a "dns" variable, and **MUST** be chosen such that the result after DoH URI Template expansion (Section 6 of [RFC8484]) is always a valid and functional ":path" value ([RFC9113], Section 8.3.1).

When using this SVCB record, the client **MUST** send any DoH requests to the HTTP origin identified by the "https" scheme, the authentication name, and the port from the "port" SvcParam (if present). HTTP requests **MUST** be directed to the resource resulting from DoH URI Template expansion of the "dohpath" value.

Clients **SHOULD NOT** query for any HTTPS RRs when using "dohpath". Instead, the SvcParams and address records associated with this SVCB record **SHOULD** be used for the HTTPS connection, with the same semantics as an HTTPS RR. However, for consistency, service operators **SHOULD** publish an equivalent HTTPS RR, especially if clients might learn about this DoH service through a different channel.

6. Limitations

This document is concerned exclusively with the DNS transport and does not affect or inform the construction or interpretation of DNS messages. For example, nothing in this document indicates whether the service is intended for use as a recursive or authoritative DNS server. Clients need to know the intended use of services based on their context.

Not all features of this specification will be applicable or effective in all contexts:

- If the authentication name is received over an insecure channel (e.g., a glue NS record), this specification cannot prevent the client from connecting to an attacker.
- Different transports might prove to be popular for different purposes (e.g., querying a recursive resolver vs. an authoritative server). Implementors are not obligated to implement all the defined transports, although doing so is beneficial for compatibility.
- Where resolution speed is a high priority, the SVCB TargetName **SHOULD** follow the convention described in [Section 10.2](#) of [SVCB], and the use of AliasMode records ([Section 2.4.2](#) of [SVCB]) is **NOT RECOMMENDED**.

7. Examples

- A resolver known as `simple.example` that supports DNS over TLS on port 853 (implicitly, as this is its default port):

```
_dns.simple.example. 7200 IN SVCB 1 simple.example. alpn=dot
```

- A DoH-only resolver at `https://doh.example/dns-query{?dns}`. (DNS over TLS is not supported.):

```
_dns.doh.example. 7200 IN SVCB 1 doh.example. (  
  alpn=h2 dohpath=/dns-query{?dns} )
```

- A resolver known as `resolver.example` that supports:
 - DoT on `resolver.example` ports 853 (implicit in record 1) and 8530 (explicit in record 2), with "resolver.example" as the Authentication Domain Name,
 - DoQ on `resolver.example` port 853 (record 1),
 - DoH at `https://resolver.example/q{?dns}` (record 1), and
 - an experimental protocol on `fooexp.resolver.example:5353` (record 3):

```
_dns.resolver.example. 7200 IN \  
  SVCB 1 resolver.example. alpn=dot,doq,h2,h3 dohpath=/q{?dns}  
  SVCB 2 resolver.example. alpn=dot port=8530  
  SVCB 3 fooexp.resolver.example. port=5353 alpn=foo foo-info=...
```

- A name server named `ns.example.` whose service configuration is published on a different domain:

```
_dns.ns.example. 7200 IN SVCB 0 _dns.ns.nic.example.
```

8. Security Considerations

8.1. Adversary on the Query Path

This section considers an adversary who can add or remove responses to the SVCB query.

During secure transport establishment, clients **MUST** authenticate the server to its authentication name, which is not influenced by the SVCB record contents. Accordingly, this document does not mandate the use of DNSSEC. This document also does not specify how clients authenticate the name (e.g., selection of roots of trust), as this procedure might vary according to the context.

8.1.1. Downgrade Attacks

This attacker cannot impersonate the secure endpoint, but it can forge a response indicating that the requested SVCB records do not exist. For a SVCB-reliant client ([SVCB], [Section 3](#)), this only results in a denial of service. However, SVCB-optional clients will generally fall back to insecure DNS in this case, exposing all DNS traffic to attacks.

8.1.2. Redirection Attacks

SVCB-reliant clients always enforce the Authentication Domain Name, but they are still subject to attacks using the transport, port number, and "dohpath" value, which are controlled by this adversary. By changing these values in the SVCB answers, the adversary can direct DNS queries for \$HOSTNAME to any port on \$HOSTNAME and any path on "https://\$HOSTNAME". If the DNS client uses shared TLS or HTTP state, the client could be correctly authenticated (e.g., using a TLS client certificate or HTTP cookie).

This behavior creates a number of possible attacks for certain server configurations. For example, if `https://$HOSTNAME/upload` accepts any POST request as a public file upload, the adversary could forge a SVCB record containing `dohpath=/upload{?dns}`. This would cause the client to upload and publish every query, resulting in unexpected storage costs for the server and privacy loss for the client. Similarly, if two DoH endpoints are available on the same origin and the service has designated one of them for use with this specification, this adversary can cause clients to use the other endpoint instead.

To mitigate redirection attacks, a client of this SVCB mapping **MUST NOT** identify or authenticate itself when performing DNS queries, except to servers that it specifically knows are not vulnerable to such attacks. If an endpoint sends an invalid response to a DNS query, the client **SHOULD NOT** send more queries to that endpoint and **MAY** log this error. Multiple DNS services **MUST NOT** share a hostname identifier ([Section 3](#)) unless they are so similar that it is safe to allow an attacker to choose which one is used.

8.2. Adversary on the Transport Path

This section considers an adversary who can modify network traffic between the client and the alternative service (identified by the TargetName).

For a SVCB-reliant client, this adversary can only cause a denial of service. However, because DNS is unencrypted by default, this adversary can execute a downgrade attack against SVCB-optional clients. Accordingly, when the use of this specification is optional, clients **SHOULD** switch to SVCB-reliant behavior if SVCB resolution succeeds. Specifications making use of this mapping **MAY** adjust this fallback behavior to suit their requirements.

9. IANA Considerations

Per [SVCB], IANA has added the following entry to the "Service Parameter Keys (SvcParamKeys)" registry.

Number	Name	Meaning	Format Reference	Change Controller	Reference
7	dohpath	DNS-over-HTTPS path template	RFC 9461	IETF	RFC 9461

Table 1

Per [Attrleaf], IANA has added the following entry to the DNS "Underscored and Globally Scoped DNS Node Names" registry:

RR Type	_NODE NAME	Reference
SVCB	_dns	RFC 9461

Table 2

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.

- [RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", RFC 6570, DOI 10.17487/RFC6570, March 2012, <<https://www.rfc-editor.org/info/rfc6570>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [RFC9113] Thomson, M., Ed. and C. Benfield, Ed., "HTTP/2", RFC 9113, DOI 10.17487/RFC9113, June 2022, <<https://www.rfc-editor.org/info/rfc9113>>.
- [SVCB] Schwartz, B., Bishop, M., and E. Nygren, "Service Binding and Parameter Specification via the DNS (DNS SVCB and HTTPS Resource Records)", RFC 9460, DOI 10.17487/RFC9460, November 2023, <<https://www.rfc-editor.org/info/rfc9460>>.

10.2. Informative References

- [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>>.
- [DNSURI] Josefsson, S., "Domain Name System Uniform Resource Identifiers", RFC 4501, DOI 10.17487/RFC4501, May 2006, <<https://www.rfc-editor.org/info/rfc4501>>.
- [FETCH] WHATWG, "Fetch Living Standard", October 2023, <<https://fetch.spec.whatwg.org/>>.
- [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>>.
- [RFC9250] Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", RFC 9250, DOI 10.17487/RFC9250, May 2022, <<https://www.rfc-editor.org/info/rfc9250>>.

Appendix A. Mapping Summary

This table serves as a non-normative summary of the DNS mapping for SVCB.

Mapped scheme	"dns"
RR type	SVCB (64)
Name prefix	_dns for port 53, else _\$PORT._dns

Required keys	alpn or equivalent
Automatically mandatory keys	port
Special behaviors	Supports all HTTPS RR SvcParamKeys
	Overrides the HTTPS RR for DoH
	Default port is per-transport
	Cleartext fallback is discouraged

Table 3

Acknowledgments

Thanks to the many reviewers and contributors, including Andrew Campling, Peter van Dijk, Paul Hoffman, Daniel Migault, Matt Nordhoff, Eric Rescorla, Andreas Schulze, and Éric Vyncke.

Author's Address

Benjamin Schwartz
Meta Platforms, Inc.
Email: ietf@bemasc.net