

DNSOP
Internet-Draft
Intended status: Standards Track
Expires: 17 October 2025

D. Eastlake
Independent
H. Song
Futurewei Technologies
15 April 2025

A Domain Name System (DNS) Service Parameter and Resource Record for
Tunneling Information
draft-eastlake-dnsop-svcb-rr-tunnel-07

Abstract

A Domain Name System (DNS) Service Binding (SVCB) Service Parameter Type and a DNS Resource Record (RR) Type are specified for storing connection tunneling / encapsulation Information in the DNS.

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 17 October 2025.

Copyright Notice

Copyright (c) 2025 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	2
1.1. Tunneling	2
1.2. Terminology	3
2. SVCB RR Service Parameter "tunnel"	4
3. TUNNEL RR Type RDATA	4
4. Use of the Specified RRs	6
5. Acknowledgements	6
6. IANA Considerations	6
7. Security Considerations	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Appendix A. Tunnel RR Type Template	8
Authors' Addresses	9

1. Introduction

The Domain Name System (DNS) is a hierarchical, distributed, highly available database with a variety of security features used for bi-directional mapping between domain names and addresses, for email routing, and for other information [RFC1034] [RFC1035] [RFC4033]. This data is formatted into resource records (RRs) whose content type and structure are indicated by the RR Type field. General familiarity with the DNS and its terminology [RFC9499] is assumed in this document.

1.1. Tunneling

It is common for there to be a requirement to use or some benefit from using a "tunnel" or encapsulation scheme when connecting to a service/host. For a reachability use case, see Section 1.3 of [RFC9012]. Typically, this involves taking a packet with a transport header addressed to the ultimate destination, adding a tunnel header to the packet, and then adding an outer transport header before transmitting the packet out of a network interface (port). The resulting packet is illustrated in Figure 1. In some cases, such as IP-in-IP, the Tunneling Header may be null.




Figure 1: Encapsulation

The addition of the Outer Transport and Tunneling Headers will lengthen packets which may result in the need for fragmentation. Some tunneling protocol support fragmentation but for those that do not, fragmentation of the Tunneled Packet before encapsulation may be required.

This document specifies a Domain Name System (DNS) Service Binding (SVCB) Service Parameter Type and a DNS Resource Record (RR) Type for storing connection tunneling / encapsulation information in the DNS. This enables the storage and retrieval of tunneling information that may be needed to connect to a remote service or host.

1.2. Terminology

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.

The following acronyms are used in this document:

DNS - Domain Name System [RFC1034][RFC1035].

IANA - Internet Assigned Numbers Authority <www.iana.org>.

RDATA - The data portion of an RR.

RR - DNS Resource Record.

RRTYPE - The type field in an RR.

SVCB - Service Binding.

2. SVCB RR Service Parameter "tunnel"

The SVCB (Service Binding) RR is specified in [RFC9460]. It provides, when used in the "Service Mode", for the encoding of a variety of "Service Parameters" to assist in connecting to a service.

The "tunnel" SVCB Service Parameter, whose numeric key value is TBD1, has a value consisting of the Tunnel Type, Tunnel Parameters Length, and Tunnel Parameters TLVs. It uses the same Tunnel Type codes and parameter TLVs as are specified for the BGP Tunnel Encapsulation Attribute in [RFC9012] as shown in Figure 2. The presentation format for this value is hexadecimal.

```

0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Tunnel Type           | Tunnel Parameters Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
/                                     /
/           Tunnel Parameters TLVs (variable length)           /
/                                                                     /
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: SVCB tunnel Service Parameter Value

For further details on the fields in Figure 2 see Section 3 which re-uses these fields with the same names and specifies those details.

3. TUNNEL RR Type RDATA

The RDATA for this RR type includes the following as further explained below and illustrated in Figure 3:

- * tunneling information in the format used in the BGP Tunnel Encapsulation Attribute [RFC9012], and
- * a domain name that maps to the Inner Transport Header destination.

The RRTYPE Code for the TUNNEL RR is TBD2.

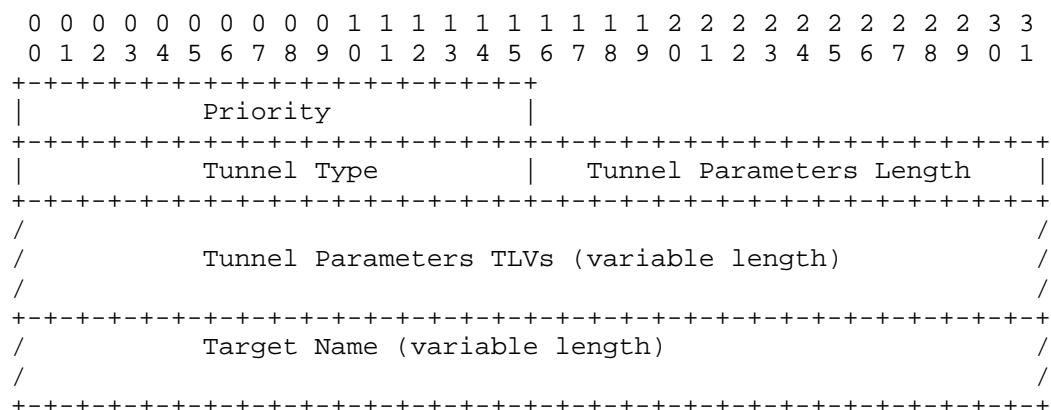


Figure 3: TUNNEL RRTYPE Data

The fields in Figure 3 are as explained below. The contiguous Tunnel Type, Tunnel Parameters Length, and Tunnel Parameters TLV (Value) as a block of data are identical to the "Tunnel Encapsulation TLV" specified in [RFC9012] ("The BGP Tunnel Encapsulation Attribute").

Priority - This field is a two-byte unsigned integer using network byte order that is the priority of using this tunnel to the target. A client **MUST** use the tunnel with the lowest priority RR that meets the following conditions:

- * The client implements the Tunnel Type.
- * The client can resolve the Target Name.
- * The type of packet being tunneled is not prohibited by an optional Protocol Type Tunnel Parameters TLV (see Section 3.4.1 of [RFC9012]). For example, the tunneling could be restricted to TCP packets.

Tunnel Type - This is the Tunnel Type from the IANA "BFP Tunnel Encapsulation Attribute Tunnel Types" registry as specified in [RFC9012].

Tunnel Parameters Length - A two-byte unsigned integer using network byte order giving the number of octets in the Tunnel Parameters TLVs field. Necessary because that field is not self-terminating.

Tunnel Parameters TLVs - This field consists of "Tunnel Encapsulation Attribute Sub-TLVs" as specified in [RFC9012]. These TLVs can specify a variety of parameters, including the following, which may be useful in constructing the Outer Transport Header (Figure 1):

- * Tunnel Egress Endpoint
- * Differentiated Services Field [RFC2474]
- * UDP Destination Port

Target Name - The uncompressed domain name of the ultimate destination in DNS wire encoding format. Used to obtain the destination address for the construction of the Inner Transport Header as shown in Figure 1.

4. Use of the Specified RRs

A client/application seeking to send packets to a host or service can query the DNS using the name of the host or service for the TUNNEL RR. If that name is found and one or more TUNNEL RRs are returned, it can use the highest priority TUNNEL RR for which it has implemented the Tunnel Type indicated in that RR to create and populate a tunneling header as shown in Figure 1. The Target Name in that TUNNEL RR can then be used to obtain an address for use in the Outer Transport Header as also shown in Figure 1. With these headers, a packet is then transmitted.

Where a client/application is already using the SVCB RR [RFC9012], similar logic applies using the tunnel SVCB Service Parameter.

5. Acknowledgements

The suggestions and comments of the following persons are gratefully acknowledged:

tbd

6. IANA Considerations

IANA is requested to assign a value from the Service Parameter Keys Registry on the "DNS Service Bindings (SVCB)" IANA web page as follows:

Number	Name	Meaning	Reference
TBD1	tunnel	Tunneling information	[this document]

Table 1

IANA is requested to assign a TUNNEL RR Type (TBD2) as in the template in Appendix A.

7. Security Considerations

For SVCB Security Considerations, see [RFC9460]. For Tunnel Encapsulation attribute Security Considerations, see [RFC9012]

More tbd

8. References

8.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [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>>.
- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9460] Schwartz, B., Bishop, M., and E. Nygren, "Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records)", RFC 9460, DOI 10.17487/RFC9460, November 2023, <<https://www.rfc-editor.org/info/rfc9460>>.

8.2. Informative References

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC9499] Hoffman, P. and K. Fujiwara, "DNS Terminology", BCP 219, RFC 9499, DOI 10.17487/RFC9499, March 2024, <<https://www.rfc-editor.org/info/rfc9499>>.

Appendix A. Tunnel RR Type Template

A. Submission Date: tbd

B.1 Submission Type: ☒ New RRTYPE ☐ Modification to RRTYPE

B.2 Kind of RR: ☒ Data RR ☐ Meta-RR

C. Contact Information for submitter (will be publicly posted):

Name: Donald Eastlake Email Address: d3e3e3@gmail.com

International telephone number: +1-508-333-2270

Other contact handles:

D. Motivation for the new RRTYPE application.

Need to store tunneling information in the DNS.

E. Description of the proposed RR type.

See draft-eastlake-dnsop-svcb-rr-tunnel

F. What existing RRTYPE or RRTYPEs come closest to filling that need and why are they unsatisfactory?

The SRV RR provides connection information for a service/host but not tunneling information.

G. What mnemonic is requested for the new RRTYPE (optional)?

TUNNEL

H. Does the requested RRTYPE make use of any existing IANA registry or require the creation of a new IANA subregistry in DNS Parameters? If so, please indicate which registry is to be used or created. If a new subregistry is needed, specify the allocation policy for it and its initial contents.

Makes use of the Border Gateway Protocol (BGP) Tunnel Encapsulation Registry and subsidiary Registries under the encapsulation registry. Does not create a new registry.

I. Does the proposal require/expect any changes in DNS servers/resolvers that prevent the new type from being processed as an unknown RRTYPE (see [RFC3597])?

No.

J. Comments: None.

Authors' Addresses

Donald Eastlake
Independent
2386 Panoramic Circle
Apopka, FL 32703
United States of America
Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

Haoyu Song
Futurewei Technologies
2220 Central Expressway
Santa Clara, CA 95050
United States of America
Email: haoyu.song@futurewei.com