

Internet Engineering Task Force (IETF)
Request for Comments: 9013
Category: Standards Track
ISSN: 2070-1721

X. Xu, Ed.
Capitalonline
B. Decraene, Ed.
Orange
R. Raszuk
NTT Network Innovations
L. Contreras
Telefonica I+D
L. Jalil
Verizon
April 2021

OSPF Advertisement of Tunnel Encapsulations

Abstract

Networks use tunnels for a variety of reasons. A large variety of tunnel types are defined, and the tunnel encapsulator router needs to select a type of tunnel that is supported by the tunnel decapsulator router. This document defines how to advertise, in OSPF Router Information Link State Advertisements (LSAs), the list of tunnel encapsulations supported by the tunnel decapsulator.

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/rfc9013>.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Terminology
3. Tunnel Encapsulations TLV
4. Tunnel Sub-TLV
5. Tunnel Parameter Sub-TLVs
 - 5.1. Encapsulation Sub-TLV
 - 5.2. Protocol Type Sub-TLV

- 5.3. Tunnel Egress Endpoint Sub-TLV
- 5.4. Color Sub-TLV
- 5.5. Load-Balancing Block Sub-TLV
- 5.6. DS Field Sub-TLV
- 5.7. UDP Destination Port Sub-TLV
- 6. Operation
- 7. IANA Considerations
 - 7.1. OSPF Router Information (RI) TLVs Registry
 - 7.2. OSPF Tunnel Parameter Sub-TLVs Registry
- 8. Security Considerations
- 9. References
 - 9.1. Normative References
 - 9.2. Informative References

Acknowledgements

Contributors

Authors' Addresses

1. Introduction

Networks use tunnels for a variety of reasons, such as:

- * Partial deployment of IPv6 in IPv4 networks or IPv4 in IPv6 networks, as described in [RFC5565], where IPvx tunnels are used between IPvx-enabled routers so as to traverse non-IPvx routers.
- * Remote Loop-Free Alternate (RLFA) repair tunnels as described in [RFC7490], where tunnels are used between the Point of Local Repair and the selected PQ node.

The tunnel encapsulator router needs to select a type of tunnel that is supported by the tunnel decapsulator router. This document defines how to advertise, in OSPF Router Information Link State Advertisements (LSAs), the list of tunnel encapsulations supported by the tunnel decapsulator. In this document, OSPF refers to both OSPFv2 [RFC2328] and OSPFv3 [RFC5340].

2. Terminology

This memo makes use of the terms defined in [RFC7770].

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. Tunnel Encapsulations TLV

Routers advertise their supported tunnel encapsulation type(s) by advertising a new TLV of the OSPF Router Information (RI) Opaque LSA [RFC7770], referred to as the "Tunnel Encapsulations TLV". This TLV is applicable to both OSPFv2 and OSPFv3.

The Type code of the Tunnel Encapsulations TLV is 13, the Length value is variable, and the Value field contains one or more Tunnel Sub-TLVs, as defined in Section 4. Each Tunnel Sub-TLV indicates a particular encapsulation format that the advertising router supports, along with the parameters corresponding to the tunnel type.

The Tunnel Encapsulations TLV MAY appear more than once within a given OSPF Router Information (RI) Opaque LSA. If the Tunnel Encapsulations TLV appears more than once in an OSPF Router Information LSA, the set of all Tunnel Sub-TLVs from all Tunnel Encapsulations TLVs SHOULD be considered. The scope of the advertisement depends on the application, but it is recommended that it SHOULD be domain wide.

4. Tunnel Sub-TLV

The Tunnel Sub-TLV is structured as shown in Figure 1.

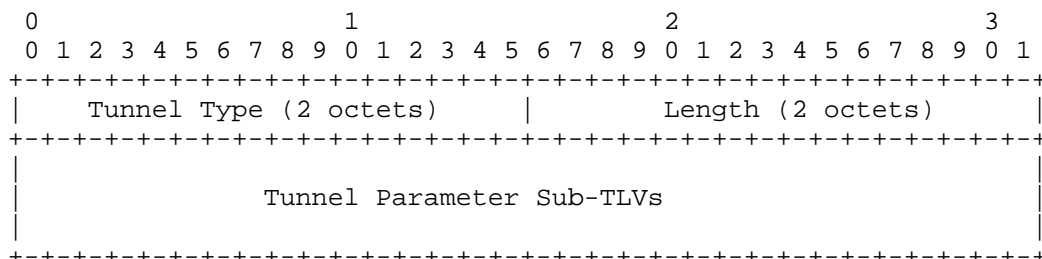


Figure 1: Tunnel Sub-TLV

Tunnel Type (2 octets): Identifies the type of tunneling technology signaled. Tunnel types are shared with the BGP extension [RFC9012] and hence are defined in the IANA registry "BGP Tunnel Encapsulation Attribute Tunnel Types". Unknown tunnel types are to be ignored upon receipt.

Length (2 octets): Unsigned 16-bit integer indicating the total number of octets of the Tunnel Parameter Sub-TLVs field.

Tunnel Parameter Sub-TLVs (variable): Zero or more Tunnel Parameter Sub-TLVs, as defined in Section 5.

If a Tunnel Sub-TLV is invalid, it MUST be ignored and skipped. However, other Tunnel Sub-TLVs MUST be considered.

5. Tunnel Parameter Sub-TLVs

A Tunnel Parameter Sub-TLV is structured as shown in Figure 2.

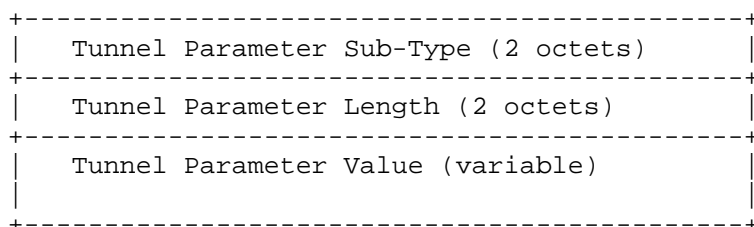


Figure 2: Tunnel Parameter Sub-TLV

Tunnel Parameter Sub-Type (2 octets): Each sub-type defines a parameter of the Tunnel Sub-TLV. Sub-types are registered in the IANA registry "OSPF Tunnel Parameter Sub-TLVs" (see Section 7.2).

Tunnel Parameter Length (2 octets): Unsigned 16-bit integer indicating the total number of octets of the Tunnel Parameter Value field.

Tunnel Parameter Value (variable): Encodings of the Value field depend on the sub-TLV type. The following subsections define the encoding in detail.

Any unknown Tunnel Parameter sub-type MUST be ignored and skipped upon receipt. When a reserved value (see Section 7.2) is seen in an LSA, it MUST be treated as an invalid Tunnel Parameter Sub-TLV. When a Tunnel Parameter Value has an incorrect syntax or semantics, it MUST be treated as an invalid Tunnel Parameter Sub-TLV. If a Tunnel Parameter Sub-TLV is invalid, its Tunnel Sub-TLV MUST be ignored. However, other Tunnel Sub-TLVs MUST be considered.

5.1. Encapsulation Sub-TLV

This sub-TLV type is 1. The syntax, semantics, and usage of its Value field are defined in Section 3.2 ("Encapsulation Sub-TLVs for Particular Tunnel Types") of [RFC9012].

5.2. Protocol Type Sub-TLV

This sub-TLV type is 2. The syntax, semantics, and usage of its Value field are defined in Section 3.4.1 ("Protocol Type Sub-TLV") of [RFC9012].

5.3. Tunnel Egress Endpoint Sub-TLV

The Tunnel Egress Endpoint Sub-TLV specifies the address of the egress endpoint of the tunnel -- that is, the address of the router that will decapsulate the payload.

This sub-TLV type is 3. It MUST be present once and only once in a given Tunnel Sub-TLV. The Value field contains two subfields:

- * a two-octet Address Family subfield
- * an Address subfield, whose length depends upon the Address Family

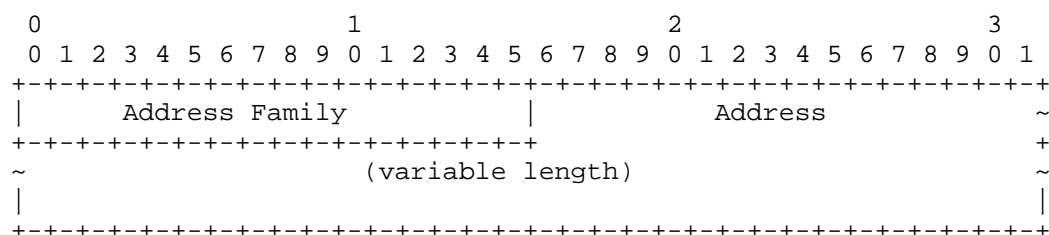


Figure 3: Tunnel Egress Endpoint Sub-TLV

The Address Family subfield contains a value from IANA's "Address Family Numbers" registry. In this document, we assume that the Address Family is either IPv4 or IPv6; use of other address families is outside the scope of this document.

If the Address Family subfield contains the value for IPv4, the Address subfield MUST contain an IPv4 address (a /32 IPv4 prefix). In this case, the Length field of the Tunnel Egress Endpoint Sub-TLV MUST contain the value 6.

If the Address Family subfield contains the value for IPv6, the address subfield MUST contain an IPv6 address (a /128 IPv6 prefix). In this case, the Length field of the Tunnel Egress Endpoint Sub-TLV MUST contain the value 18 (0x12). IPv6 link-local addresses are not valid values of the IP address field.

5.4. Color Sub-TLV

This sub-TLV type is 4. It may appear zero or more times in a given Tunnel Sub-TLV. The Value field is a 4-octet opaque unsigned integer.

The color value is user-defined and configured locally on the advertising routers. It may be used by service providers to define policies on the tunnel encapsulator routers, for example, to control the selection of the tunnel to use.

This color value can be referenced by BGP routes carrying the Color

Extended Community [RFC9012]. If the tunnel is used to reach the BGP next hop of BGP routes, then attaching a Color Extended Community to those routes expresses the willingness of the BGP speaker to use a tunnel of the same color.

5.5. Load-Balancing Block Sub-TLV

This sub-TLV type is 5. The syntax, semantics, and usage of its Value field are defined in [RFC5640].

5.6. DS Field Sub-TLV

This sub-TLV type is 6. The syntax, semantics, and usage of its Value field are defined in Section 3.3.1 ("DS Field") of [RFC9012].

5.7. UDP Destination Port Sub-TLV

This sub-TLV type is 7. The syntax, semantics, and usage of its Value field are defined in Section 3.3.2 ("UDP Destination Port") of [RFC9012].

6. Operation

The advertisement of a Tunnel Encapsulations Sub-TLV indicates that the advertising router supports a particular tunnel decapsulation along with the parameters to be used for the tunnel. The decision to use that tunnel is driven by the capability of the tunnel encapsulator router to support the encapsulation type and the policy on the tunnel encapsulator router. The Color Sub-TLV (see Section 5.4) may be used as an input to this policy. Note that some tunnel types may require the execution of an explicit tunnel setup protocol before they can be used to transit data.

A tunnel MUST NOT be used if there is no route toward the IP address specified in the Tunnel Egress Endpoint Sub-TLV (see Section 5.3) or if the route is not advertised in the same OSPF domain.

7. IANA Considerations

7.1. OSPF Router Information (RI) TLVs Registry

IANA has allocated the following new code point in the "OSPF Router Information (RI) TLVs" registry.

Value	TLV Name	Reference
13	Tunnel Encapsulations	RFC 9013

Table 1: Addition to OSPF Router Information (RI) TLVs Registry

7.2. OSPF Tunnel Parameter Sub-TLVs Registry

IANA has created a new subregistry called the "OSPF Tunnel Parameter Sub-TLVs" registry under the "Open Shortest Path First (OSPF) Parameters" registry. The registration procedures are as follows:

- * The values in the range 1-34999 are to be allocated using the "Standards Action" registration procedure defined in [RFC8126].
- * The values in the range 35000-65499 are to be allocated using the "First Come First Served" registration procedure.

The initial contents of the registry are as follows:

Value	TLV Name	Reference
0	Reserved	RFC 9013
1	Encapsulation	RFC 9013 & RFC 9012
2	Protocol Type	RFC 9013 & RFC 9012
3	Endpoint	RFC 9013
4	Color	RFC 9013
5	Load-Balancing Block	RFC 9013 & RFC 5640
6	DS Field	RFC 9013 & RFC 9012
7	UDP Destination Port	RFC 9013 & RFC 9012
8-65499	Unassigned	
65500-65534	Experimental	RFC 9013
65535	Reserved	RFC 9013

Table 2: Initial Contents of OSPF Tunnel Parameter Sub-TLVs Registry

8. Security Considerations

Security considerations applicable to softwires can be found in the mesh framework [RFC5565]. In general, security issues of the tunnel protocols signaled through this OSPF capability extension are inherited.

If a third party is able to modify any of the information that is used to form encapsulation headers, choose a tunnel type, or choose a particular tunnel for a particular payload type, user data packets may end up getting misrouted, misdelivered, and/or dropped. However, since an OSPF routing domain is usually a well-controlled network under a single administrative domain, the possibility of the above attack is very low.

We note that the last paragraph of Section 6 forbids the establishment of a tunnel toward arbitrary destinations. It prohibits a destination outside of the OSPF domain. This prevents a third party that has gained access to an OSPF router from being able to send the traffic to other destinations, e.g., for inspection purposes.

Security considerations for the base OSPF protocol are covered in [RFC2328] and [RFC5340].

9. References

9.1. 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>>.
- [RFC5640] Filsfils, C., Mohapatra, P., and C. Pignataro, "Load-Balancing for Mesh Softwires", RFC 5640,

DOI 10.17487/RFC5640, August 2009,
<<https://www.rfc-editor.org/info/rfc5640>>.

- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <<https://www.rfc-editor.org/info/rfc7770>>.
- [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>>.
- [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>>.

9.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", RFC 5512, DOI 10.17487/RFC5512, April 2009, <<https://www.rfc-editor.org/info/rfc5512>>.
- [RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", RFC 5565, DOI 10.17487/RFC5565, June 2009, <<https://www.rfc-editor.org/info/rfc5565>>.
- [RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.

Acknowledgements

This document is partially inspired by [RFC5512].

The authors would like to thank Greg Mirsky, John E. Drake, Carlos Pignataro, and Karsten Thomann for their valuable comments on this document. Special thanks should be given to Acee Lindem for his multiple detailed reviews of this document and help. The authors would like to thank Pete Resnick, Joe Touch, David Mandelberg, Sabrina Tanamal, Tim Wicinski, and Amanda Baber for their Last Call reviews. The authors also thank Spencer Dawkins, Mirja Khlewind, Ben Campbell, Benoit Claise, Alvaro Retana, Adam Roach, and Suresh Krishnan for their AD reviews.

Contributors

Uma Chunduri
Huawei

Email: uma.chunduri@gmail.com

Authors' Addresses

Xiaohu Xu (editor)
Capitalonline

Email: xiaohu.xu@capitalonline.net

Bruno Decraene (editor)
Orange

Email: bruno.decraene@orange.com

Robert Raszuk
NTT Network Innovations
940 Stewart Dr
Sunnyvale, CA 94085
United States of America

Email: robert@raszuk.net

Luis M. Contreras
Telefonica I+D

Email: luismiguel.contrerasmurillo@telefonica.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com