

OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2018

X. Xu, Ed.
Huawei
B. Decraene, Ed.
Orange
R. Raszuk
Bloomberg LP
L. Contreras
Telefonica I+D
L. Jalil
Verizon
October 9, 2017

**The Tunnel Encapsulations OSPF Router Information
draft-ietf-ospf-encapsulation-cap-09**

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 which is supported by the tunnel decapsulator router. This document defines how to advertise, in OSPF Router Information Link State Advertisement (LSAs), the list of tunnel encapsulations supported by the tunnel decapsulator.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 April 12, 2018.

Copyright Notice

Copyright (c) 2017 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](#)
- [2.](#) Terminology [3](#)
- [3.](#) Tunnel Encapsulations TLV [3](#)
- [4.](#) Tunnel Sub-TLV [3](#)
- [5.](#) Tunnel Parameter Sub-TLVs [4](#)
 - [5.1.](#) Encapsulation Sub-TLV [5](#)
 - [5.2.](#) Protocol Type Sub-TLV [5](#)
 - [5.3.](#) Endpoint Sub-TLV [5](#)
 - [5.4.](#) Color Sub-TLV [6](#)
 - [5.5.](#) Load-Balancing Block Sub-TLV [6](#)
 - [5.6.](#) IP QoS Field [6](#)
 - [5.7.](#) UDP Destination Port [7](#)
- [6.](#) Operation [7](#)
- [7.](#) IANA Considerations [7](#)
 - [7.1.](#) OSPF Router Information [7](#)
 - [7.2.](#) Tunnel Parameter Sub-TLVs Registry [7](#)
- [8.](#) Security Considerations [8](#)
- [9.](#) Contributors [9](#)
- [10.](#) Acknowledgements [9](#)
- [11.](#) References [9](#)
 - [11.1.](#) Normative References [9](#)
 - [11.2.](#) Informative References [10](#)
- Authors' Addresses [10](#)

1. Introduction

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

- o 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.

- o 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 which is supported by the tunnel decapsulator router. This document defines how to advertise, in OSPF Router Information Link State Advertisement (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](#)].

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 is TBD1, 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 TLV 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 follows:

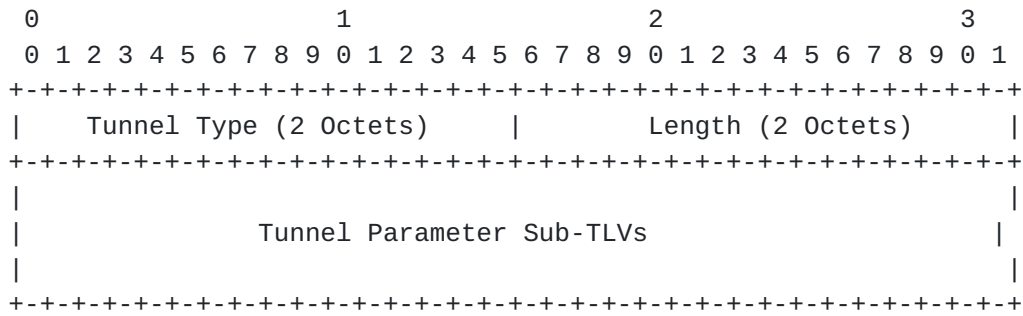


Figure 1: Tunnel Sub-TLV

Tunnel Type (2 octets): Identifies the type of tunneling technology signaled. Tunnel types are shared with the BGP extension [[I-D.ietf-idr-tunnel-encaps](#)] 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 value field.

Value (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 follows:




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" [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 as enumerated above. The following sub-sections 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 semantic, 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, semantic, and usage of its value field are defined in [Section 3.2](#) "Encapsulation Sub-TLVs for Particular Tunnel Types" of [[I-D.ietf-idr-tunnel-encaps](#)].

5.2. Protocol Type Sub-TLV

This Sub-TLV type is 2. The syntax, semantic, and usage of its value field are defined in [Section 3.4.1](#) "Protocol Type sub-TLV" of [[I-D.ietf-idr-tunnel-encaps](#)].

5.3. Endpoint Sub-TLV

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

a two-octet Address Family sub-field

an Address sub-field, whose length depends upon the Address Family.

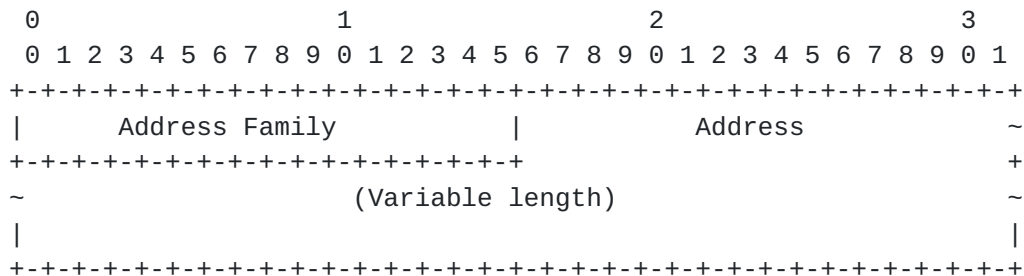


Figure 3: Endpoint Sub-TLV

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 Remote Endpoint sub-TLV MUST contain the value 6.

If the Address Family subfield contains the value for IPv6, the address sub-field MUST contain an IPv6 address (a /128 IPv6 prefix). In this case, the length field of Remote 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 Color Extended Community [[I-D.ietf-idr-tunnel-encaps](#)]. 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, semantic, and usage of its value field are defined in [[RFC5640](#)].

5.6. IP QoS Field

This Sub-TLV type is 6. The syntax, semantic, and usage of its value field are defined in [Section 3.3.1](#) "IPv4 DS Field" of [[I-D.ietf-idr-tunnel-encaps](#)].

5.7. UDP Destination Port

This Sub-TLV type is 7. The syntax, semantic, and usage of its value field are defined in [Section 3.3.2](#) "UDP Destination Port" of [[I-D.ietf-idr-tunnel-encaps](#)].

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 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

This document requests IANA to allocate a new code point from the OSPF Router Information (RI) registry.

Value	TLV Name	Reference
-----	-----	-----
TBD1	Tunnel Encapsulations	This document

Figure 4: Tunnel Encapsulation Router Information

7.2. Tunnel Parameter Sub-TLVs Registry

This document requests IANA to create, under "Open Shortest Path First (OSPF) Parameters", a new registry "OSPF Tunnel Parameter Sub-TLVs" with the following registration procedure:

The values in the range 1-34999 are to be allocated using the "Standards Action" registration procedure as defined in [[RFC8126](#)].

The values in the range 35000-65499 are to be allocated using the "First Come, First Served" registration procedure.

Registry Name: OSPF Tunnel Parameter Sub-TLVs

Value	Name	Reference
0	Reserved	This document
1	Encapsulation	This document & [I-D.ietf-idr-tunnel-encaps]
2	Protocol Type	This document & [I-D.ietf-idr-tunnel-encaps]
3	Endpoint	This document
4	Color	This document
5	Load-Balancing Block	This document & [RFC5640]
6	IP QoS	This document & [I-D.ietf-idr-tunnel-encaps]
7	UDP Destination Port	This document & [I-D.ietf-idr-tunnel-encaps]
8-65499	Unassigned	
65500-65534	Experimental	This document
65535	Reserved	This document

Figure 5: 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, to choose a tunnel type, or to choose a particular tunnel for a particular payload type, user data packets may end up getting misrouted, mis-delivered, 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](#) forbid the establishment of a tunnel toward arbitrary destinations. It prohibits a destination outside of the OSPF domain. This avoid that a third-party gaining access to an OSPF router be 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. Contributors

Uma Chunduri
Huawei
Email: uma.chunduri@gmail.com

10. 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, Amanda Baber for their Last Call reviews and thank Spencer Dawkins, Mirja Kuehlewind, Ben Campbell, Benoit Claise, Alvaro Retana, Adam Roach and Suresh Krishnan for their AD reviews.

11. References

11.1. Normative References

- [I-D.ietf-idr-tunnel-encaps]
Rosen, E., Patel, K., and G. Velde, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-07](#) (work in progress), July 2017.
- [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>>.

11.2. 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>>.

Authors' Addresses

Xiaohu Xu (editor)
Huawei

Email: xuxiaohu@huawei.com

Bruno Decraene (editor)
Orange

Email: bruno.decraene@orange.com

Robert Raszuk
Bloomberg LP

Email: robert@raszuk.net

Luis M. Contreras
Telefonica I+D

Email: luismiguel.contrerasmurillo@telefonica.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com