OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2018

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

September 10, 2017

Advertising Tunnel Encapsulation Capabilities in OSPF
draft-ietf-ospf-encapsulation-cap-07

Abstract

Networks use tunnels for a variety of reasons. A large variety of tunnel types are defined and the ingress tunnel router needs to select a type of tunnel which is supported by the egress tunnel router and itself. This document defines how to advertise the tunnel encapsulation capabilities of egress tunnel routers in OSPF Router Information Link State Advertisement (LSAs).

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

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 March 14, 2018.

Table of Contents

1. Introduction 2
2. Terminology 3
3. Tunnel Encapsulation Capabilities TLV 3
4. Tunnel Encapsulation Type Sub-TLVs 3
5. Tunnel Encapsulation Attribute 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 5
5.5. Load-Balancing Block Sub-TLV 6
5.6. IP QoS Field 6
5.7. UDP Destination Port 6
6. IANA Considerations 6
6.1. OSPF Router Information 6
6.2. Tunnel Encapsulation Attribute Sub-TLVs Registry 6
7. Security Considerations 7
8. Contributors 7
9. Acknowledgements 7
10. References 8
10.1. Normative References 8
10.2. Informative References 8
Authors' Addresses 9

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 ingress tunnel router needs to select a type of tunnel which is supported by the egress tunnel router and itself. This document describes how to use OSPF Router Information Link State Advertisements (LSAs) to advertise the tunneling capabilities of OSPF routers acting as egress tunnel routers. 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 Encapsulation Capabilities TLV

Routers advertise their supported encapsulation type(s) by advertising a new TLV of the OSPF Router Information (RI) Opaque LSA [RFC7770], referred to as the Tunnel Encapsulation Capabilities TLV. This TLV is applicable to both OSPFv2 and OSPFv3. The Tunnel Encapsulation Capabilities TLV SHOULD NOT appear more than once within a given OSPF Router Information (RI) Opaque LSA. If the Tunnel Encapsulation Capabilities TLV appears more than once in an OSPF Router Information LSA, only the first occurrence MUST be processed and others SHOULD be ignored. The scope of the advertisement depends on the application but it is recommended that it SHOULD be domain-wide. The Type code of the Tunnel Encapsulation Capabilities TLV is TBD1, the Length value is variable, and the Value field contains one or more Tunnel Encapsulation Type Sub-TLVs (see Section 4). Each Encapsulation Type Sub-TLVs indicates a particular encapsulation format that the advertising router supports along with the parameters to be used for the tunnel.

4. Tunnel Encapsulation Type Sub-TLVs

The Tunnel Encapsulation Type Sub-TLV is structured as follows:

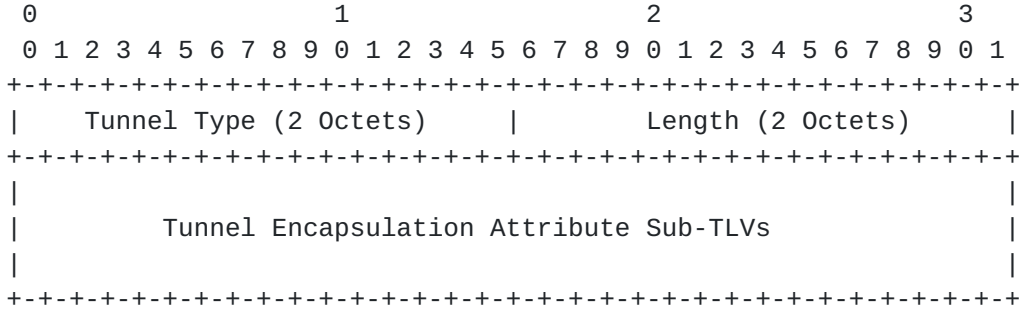

Tunnel Type (2 octets): Identifies the type of tunneling technology being signaled. Tunnel types are shared with the BGP extension [I-D.ietf-idr-tunnel-encaps] and hence are defined in the IANA registry "BGP Tunnel Encapsulation Attribute Tunnel Types". Unknown types are to be ignored and skipped upon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

Length (2 octets): Unsigned 16-bit integer indicating the total number of octets of the value field. Note that this is a padding to be ignored if the length field is longer than the field indicated by the sub-TLVs.

Value (variable): Zero or more Tunnel Encapsulation Attribute Sub-TLVs as defined in Section 5.

5. Tunnel Encapsulation Attribute Sub-TLVs

Tunnel Encapsulation Attribute Sub-TLV are structured as follows:

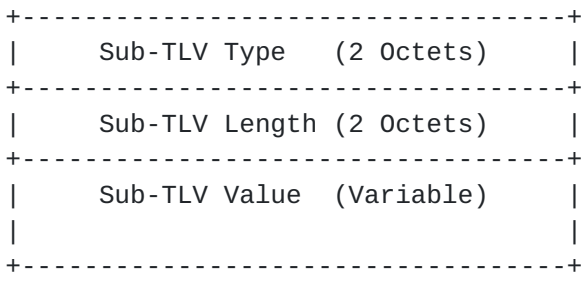

Sub-TLV Type (2 octets): Each Sub-TLV type defines a certain property of the tunnel TLV that contains this Sub-TLV. Types are registered in the IANA registry "OSPF Tunnel Encapsulation Attribute Sub-TLVs" Section 6.2.

Sub-TLV Length (2 octets): Unsigned 16-bit integer indicating the total number of octets of the Sub-TLV value field.

Sub-TLV 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 Sub-TLVs MUST be deemed as invalid Sub-TLVs and therefore MUST be ignored and skipped upon receipt. When a reserved value (See Section 6.2) is seen in an LSA, it SHOULD be treated as an invalid Sub-TLV. If a Sub-TLV is invalid, its Tunnel Encapsulation Type TLV MUST be ignored and skipped. However, other Tunnel Encapsulation Type TLVs MUST be considered.

The advertisement of an Encapsulation Type Sub-TLV (See Section 5.1) indicates that the advertising router support a particular tunnel

encapsulation along with the parameters to be used for the tunnel. The decision to use that tunnel is driven by the capability of the ingress router to support the encapsulation type and the policy on the ingress 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 carry 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 by the router advertising the Tunnel Encapsulation Attribute Sub-TLVs for the tunnel.

5.1. Encapsulation Sub-TLV

This Sub-TLV of type 1 is defined in Section 3.2 "Encapsulation Sub-TLVs for Particular Tunnel Types" of [I-D.ietf-idr-tunnel-encaps] from both a syntax and semantic standpoint.

5.2. Protocol Type Sub-TLV

This Sub-TLV of type 2 is defined in Section 3.4.1 "Protocol Type sub-TLV" of [I-D.ietf-idr-tunnel-encaps] from a syntactic, semantic, and usage standpoint.

5.3. Endpoint Sub-TLV

Type is 3. The value field carries the Network Address to be used as tunnel destination address.

If length is 4, the tunnel endpoint is an IPv4 address.

If length is 16, the tunnel endpoint is an IPv6 address.

5.4. Color Sub-TLV

Type is 4. 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 ingress tunnel 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 attached to those routes express the willingness of the BGP speaker to use a tunnel of the same color.

5.5. Load-Balancing Block Sub-TLV

This Sub-TLV of type 5 is defined in [RFC5640] from a syntactic, semantic and usage standpoint.

5.6. IP QoS Field

This Sub-TLV of type 6 is defined in Section 3.3.1 "IPv4 DS Field" of [I-D.ietf-idr-tunnel-encaps] from a syntactic, semantic and usage standpoint.

5.7. UDP Destination Port

This Sub-TLV of type 7 is defined in Section 3.3.2 "UDP Destination Port" of [I-D.ietf-idr-tunnel-encaps] from a syntactic, semantic and usage standpoint.

6. IANA Considerations

6.1. 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 Encapsulation Capabilities	This document

6.2. Tunnel Encapsulation Attribute Sub-TLVs Registry

This document requests IANA to create a new registry "Tunnel Encapsulation Attribute Sub-TLVs" with the following registration procedure:

The values in the range 1-255 are to be allocated using the "Standards Action" registration procedure as defined in [RFC5226].

The values in the range 256-65499 are to be allocated using the "First Come, First Served" registration procedure.

Registry Name: OSPF Tunnel Encapsulation Attribute 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

7. 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, misdelivered, and/or dropped. However, since an OSPF routing domain is usually well-controlled and well-managed network, the possiblity of the above risk is very low.

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

8. Contributors

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

9. 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.

10. References

10.1. 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>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <https://www.rfc-editor.org/info/rfc5226>.

[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>.

10.2. 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