

OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 6, 2020

X. Xu
Alibaba Inc
S. Kini

P. Psenak
C. Filsfils
Cisco
S. Litkowski
Orange
September 3, 2019

**Signaling Entropy Label Capability and Entropy Readable Label-stack
Depth Using OSPF
draft-ietf-ospf-mpls-elc-09**

Abstract

Multiprotocol Label Switching (MPLS) has defined a mechanism to load-balance traffic flows using Entropy Labels (EL). An ingress Label Switching Router (LSR) cannot insert ELs for packets going into a given tunnel unless an egress LSR has indicated via signaling that it has the capability to process ELs, referred to as Entropy Label Capability (ELC), on that tunnel. In addition, it would be useful for ingress LSRs to know each LSR's capability of reading the maximum label stack depth and performing EL-based load-balancing, referred to as Entropy Readable Label Depth (ERLD). This document defines a mechanism to signal these two capabilities using OSPF and OSPFv3. These mechanism is particularly useful in the environment where Segment Routing (SR) is used, where label advertisements are done via protocols like OSPF and OSPFv3.

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 March 6, 2020.

Copyright Notice

Copyright (c) 2019 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.	Advertising ELC Using OSPF	3
3.1.	Advertising ELC Using OSPFv2	4
3.2.	Advertising ELC Using OSPFv3	4
4.	Advertising ERLD Using OSPF	4
5.	Acknowledgements	4
6.	BGP-LS Extension	4
7.	IANA Considerations	5
8.	Security Considerations	5
9.	References	5
9.1.	Normative References	5
9.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

[RFC6790] describes a method to load-balance Multiprotocol Label Switching (MPLS) traffic flows using Entropy Labels (EL). It also introduces the concept of Entropy Label Capability (ELC) and defines the signaling of this capability via MPLS signaling protocols. Recently, mechanisms have been defined to signal labels via link-state Interior Gateway Protocols (IGP) such as OSPF [[I-D.ietf-ospf-segment-routing-extensions](#)]. In such scenarios, the

signaling mechanisms defined in [[RFC6790](#)] are inadequate. This draft defines a mechanism to signal the ELC using OSPF. This mechanism is useful when the label advertisement is also done via OSPF.

In addition, in the cases where stacked LSPs are used for whatever reasons (e.g., SR-MPLS [[I-D.ietf-spring-segment-routing-mpls](#)]), it would be useful for ingress LSRs to know each intermediate LSR's capability of reading the maximum label stack depth and performing EL-based load-balancing. This capability, referred to as Entropy Readable Label Depth (ERLD) as defined in [[I-D.ietf-mpls-spring-entropy-label](#)] may be used by ingress LSRs to determine the position of the EL label in the stack, and whether it's necessary to insert multiple ELs at different positions in the label stack.

2. Terminology

This document makes use of the terms defined in [[RFC6790](#)] and [[RFC7770](#)].

3. Advertising ELC Using OSPF

Even though ELC is a property of the node, in some cases it is advantageous to associate and advertise the ELC with the prefix. In multi-area networks, routers may not know the identity of the prefix originator in a remote area, or may not know the capabilities of such originator. Similarly, in a multi domain network, the identity of the prefix originator and its capabilities may not be known to the ingress LSR.

If a router has multiple line cards, the router **MUST NOT** announce ELC unless all of its line-cards are capable of processing ELs.

If the router supports ELs on all of its line cards, it **SHOULD** advertise the ELC with every local host prefix it advertises in OSPF.

When an OSPF Area Border Router (ABR) advertises the prefix to the connected area based on the intra-area or inter-area prefix that is reachable in some other area, it **MUST** preserve the ELC signalling for such prefix.

When an OSPF Autonomous System Boundary Router (ASBR) redistributes the prefix from another instance of the OSPF or from some other protocol, it **SHOULD** preserve the ELC signaling for the prefix. The exact mechanism used to exchange ELC between protocol instances on the ASBR is outside of the scope of this document and is implementation specific.

3.1. Advertising ELC Using OSPFv2

[RFC7684] defines the OSPFv2 Extended Prefix TLV to advertise additional attributes associated with a prefix. The OSPFv2 Extended Prefix TLV includes a one octet Flags field. A new flag in the Flags field is used to signal the ELC for the prefix:

0x20 - E-Flag (ELC Flag): Set by the advertising router to indicate that the prefix originator is capable of processing ELs.

3.2. Advertising ELC Using OSPFv3

[RFC5340] defines the OSPFv3 PrefixOptions that are advertised along with the prefix. A new bit in the OSPFv3 PrefixOptions is used to signal the ELC for the prefix:

0x04 - E-Flag (ELC Flag): Set by the advertising router to indicate that the prefix originator is capable of processing ELs.

4. Advertising ERLD Using OSPF

A new MSD-type of the Node MSD sub-TLV

[[I-D.ietf-isis-segment-routing-msd](#)], called ERLD is defined to advertise the ERLD of a given router. The scope of the advertisement depends on the application.

Assignment of a MSD-Type for ERLD is defined in [[I-D.ietf-isis-mpls-elc](#)].

If a router has multiple line-cards with different capabilities for reading the maximum label stack depth, the router MUST advertise the smallest one.

5. Acknowledgements

The authors would like to thank Yimin Shen, George Swallow, Acee Lindem, Les Ginsberg, Ketan Talaulikar, Jeff Tantsura, Bruno Decraene and Carlos Pignataro for their valuable comments.

6. BGP-LS Extension

The OSPF extensions defined in this document can be advertised via BGP-LS [[RFC7752](#)] using existing BGP-LS TLVs.

The ELC Flag included in the OSPFv2 Extended Prefix TLV and the OSPFv3 PrefixOptions, as defined in [Section 3](#), is advertised using the Prefix Attribute Flags TLV (TLV 1170) of the BGP-LS IPv4/IPv6

Prefix NLRI Attribute as defined in section 2.3.2 of [\[I-D.ietf-idr-bgp-ls-segment-routing-ext\]](#).

The ERLD MSD-type introduced for OSPF in [Section 4](#) is advertised using the Node MSD TLV (TLV 266) of the BGP-LS Node NLRI Attribute as defined in section 3 of [\[I-D.ietf-idr-bgp-ls-segment-routing-ext\]](#).

7. IANA Considerations

This document requests IANA to allocate one flag from the OSPFv2 Extended Prefix TLV Flags registry:

0x20 - E-Flag (ELC Flag)

This document requests IANA to allocate one flag from the OSPFv3 Prefix Options registry:

0x04 - E-Flag (ELC Flag)

8. Security Considerations

The security considerations as described in [\[RFC7770\]](#) and [\[I-D.ietf-mpls-spring-entropy-label\]](#) are applicable to this document.

Incorrectly setting the E flag (ELC capable) (during origination, inter-area advertisement or redistribution) may lead to black-holing of the traffic on the egress node.

9. References

9.1. Normative References

[I-D.ietf-idr-bgp-ls-segment-routing-ext]

Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", [draft-ietf-idr-bgp-ls-segment-routing-ext-16](#) (work in progress), June 2019.

[I-D.ietf-isis-mpls-elc]

Xu, X., Kini, S., Psenak, P., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS", [draft-ietf-isis-mpls-elc-07](#) (work in progress), May 2019.

[I-D.ietf-isis-segment-routing-msd]

Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling MSD (Maximum SID Depth) using IS-IS", [draft-ietf-isis-segment-routing-msd-19](#) (work in progress), October 2018.

[I-D.ietf-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-27](#) (work in progress), December 2018.

[I-D.ietf-spring-segment-routing-mpls]

Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-22](#) (work in progress), May 2019.

[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>>.

[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>>.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", [RFC 6790](#), DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.

[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>>.

9.2. Informative References

[I-D.ietf-mpls-spring-entropy-label]
Kini, S., Kompella, K., Sivabalan, S., Litkowski, S.,
Shakir, R., and J. Tantsura, "Entropy label for SPRING
tunnels", [draft-ietf-mpls-spring-entropy-label-12](#) (work in
progress), July 2018.

Authors' Addresses

Xiaohu Xu
Alibaba Inc

Email: xiaohu.xxh@alibaba-inc.com

Sriganesh Kini

Email: sriganeshkini@gmail.com

Peter Psenak
Cisco

Email: ppsenak@cisco.com

Clarence Filsfils
Cisco

Email: cfilsfil@cisco.com

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com

