OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 21, 2017

X. Xu
Huawei
S. Kini

S. Sivabalan
C. Filsfils
Cisco
S. Litkowski
Orange

October 18, 2016

Signaling Entropy Label Capability Using OSPF
draft-ietf-ospf-mpls-elc-03

Abstract

Multi Protocol Label Switching (MPLS) has defined a mechanism to load balance traffic flows using Entropy Labels (EL). An ingress LSR cannot insert ELs for packets going into a given tunnel unless an egress LSR has indicated via signaling that it can process ELs on that tunnel. This draft defines a mechanism to signal that capability using OSPF. This mechanism is useful when the label advertisement is also done via OSPF.

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 http://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 21, 2017.

Table of Contents

1. Introduction

Multi Protocol Label Switching (MPLS) has defined a method in [RFC6790] to load balance traffic flows using Entropy Labels (EL). An ingress LSR cannot insert ELs for packets going into a given tunnel unless an egress LSR has indicated that it can process ELs on that tunnel. [RFC6790] defines the signaling of this capability (a.k.a., Entropy Label Capability - ELC) via signaling protocols. Recently, mechanisms are being defined to signal labels via link state Interior Gateway Protocols (IGP) such as OSPF [I-D.ietf-ospf-segment-routing-extensions] . In such scenario 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., SPRING-MPLS [I-D.ietf-spring-segment-routing-mpls]), it would be useful for ingress LSRs to know each LSR's capability of

reading the maximum label stack deepth. This capability, referred to as Readable Label Deepth Capability (RLDC) can be used by ingress LSRs to determine whether it's necessary to insert an EL for a given LSP tunnel in the case where there has already been at least one EL in the label stack [I-D.ietf-mpls-spring-entropy-label] . Of course, even it has been determined that it's neccessary to insert an EL for a given LSP tunnel, if the egress LSR of that LSP tunnel has not yet indicated that it can process ELs for that tunnel, the ingress LSR MUST NOT include an entropy label for that tunnel as well.

2. Terminology

This memo makes use of the terms defined in [RFC6790] and [RFC7770].

3. Non-OSPF Functional Capabilities TLV

This document defines the Router Non-OSPF Functional Capabilities TLV for advertisement in the OSPF Router Information LSA. An OSPF router advertising an OSPF RI LSA MAY include the Router Non-OSPF Functional Capabilities TLV. If included, it MUST be included in the first instance of the LSA. Additionally, the TLV MUST reflect the advertising OSPF router's actual non-OSPF functional capabilities in the flooding scope of the containing OSPF RI LSA.

The format of the Router Non-OSPF Functional Capabilities TLV is as follows:

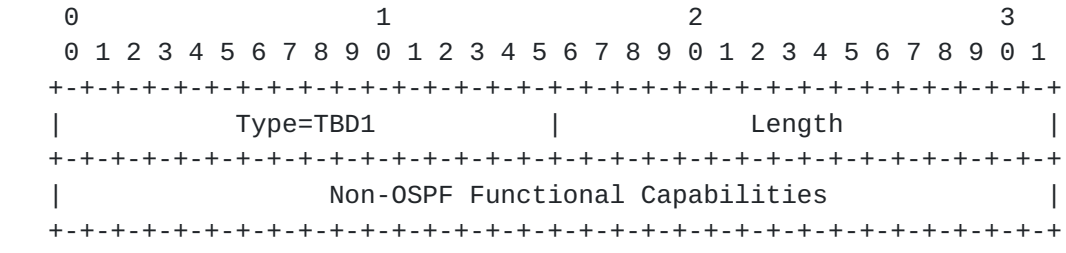

Figure 1: Non-OSPF Functional Capabilities TLV Format

Type: TBD1.

Length: indicates the length of the value portion in octets and will be a multiple of 4 octets dependent on the number of capabilities advertised. Initially, the length will be 4, denoting 4 octets of non-OSPF functional capability bits.

Value: A variable-length sequence of capability bits rounded to a multiple of 4 octets padded with undefined bits. Initially, there are 4 octets of capability bits. Bits are numbered left to right starting with the most significant bit being bit 0.

The Non-OSPF Functional Capabilities TLV MAY be followed by optional TLVs that further specify a non-OSPF functional capability. In contrast to the OSPF Router Functional Capabilities TLV, the non-OSPF functional capabilities advertised in this TLV have no impact on the OSPF protocol operation. The specifications for non-OSPF functional capabilities advertised in this TLV MUST describe protocol behavior and address backwards compatibility.

4. Advertising ELC Using OSPF

One bit of the Non-OSPF Functional Capability Bits is to be assigned by the IANA for the ELC. If a router has multiple linecards, the router MUST NOT announce the ELC unless all of its linecards are capable of processing ELs.

5. Advertising RLDC Using OSPF

A new TLV within the body of the OSPF RI LSA, called RLDC TLV is defined to advertise the capability of the router to read the maximum label stack depth. As showed in Figure 2, it is formatted as described in Section 2.3 of [RFC7770] with a Type code to be assigned by IANA and a Length of one. The Value field is set to the maximum readable label stack depth in the range between 1 to 255. The scope of the advertisement depends on the application but it is RECOMMENDED that it SHOULD be domain-wide. If a router has multiple linecards with different capabilities of reading the maximum label stack deepth, the router MUST advertise the smallest one in the RLDC TLV. This TLV is applicable to both OSPFv2 and OSPFv3.

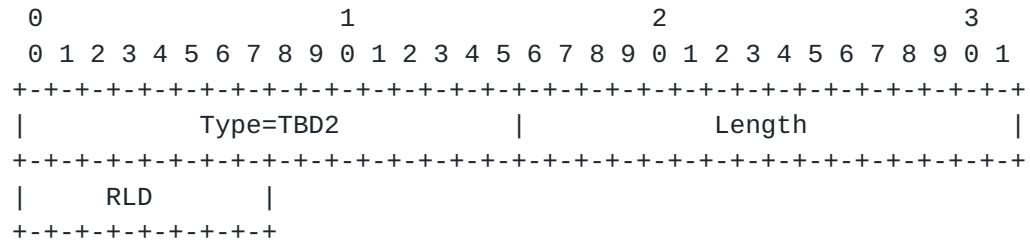

Figure 2: RLDC TLV Format

6. Usage and Applicability

The ELC is used by ingress LSRs to determine whether an EL could be inserted into a given LSP tunnel. The RLDC is used by ingress LSRs to determine whether it's necessary to insert an EL for a given LSP tunnel in the case where there has already been at least one EL in the label stack. This document only describes how to signal the ELC and RLDC using OSPF. As for how to apply those capabilities when inserting EL(s) into LSP tunnel(s), it's outside the scope of this

document and accordingly would be described in [I-D.ietf-mpls-spring-entropy-label].

7. Acknowledgements

The authors would like to thank Yimin Shen, George Swallow, Acee Lindem and Carlos Pignataro for their valuable comments.

8. IANA Considerations

This document requests IANA to allocate one TLV type from the OSPF RI TLVs registry for the Non-OSPF Functional Capabilities TLV. Futhermore, this document requests IANA to creat a subregistry for "Non-OSPF Functional Capability Bits" within the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. This subregistry is comprised of the fields Bit Number, Capability Name, and Reference. Initially, one bit is reqested to be assigned for the ELC. All Non-OSPF Functional Capability TLV additions are to be assigned through IETF Review [RFC5226].

This document also requests IANA to allocate one TLV type from the OSPF RI TLVs registry for the RLDC TLV.

9. Security Considerations

The security considerations as described in [RFC7770] is appliable to this document. This document does not introduce any new security risk.

10. References

10.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[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, <http://www.rfc-editor.org/info/rfc7770>.

10.2. Informative References

[I-D.ietf-mpls-spring-entropy-label]
Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and j. jefftant@gmail.com, "Entropy labels for source routed tunnels with label stacks", draft-ietf-mpls-spring-entropy-label-04 (work in progress), July 2016.

[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-09 (work in progress), July 2016.

[I-D.ietf-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Shakir, R., jefftant@gmail.com, j., and E. Crabbe, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-05 (work in progress), July 2016.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[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, <http://www.rfc-editor.org/info/rfc6790>.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Sriganesh Kini

Email: sriganeshkini@gmail.com

Siva Sivabalan
Cisco

Email: msiva@cisco.com

Clarence Filsfils
Cisco

Email: cfilsfil@cisco.com

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com