

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 29, 2019

X. Xu
Alibaba Inc
S. Kini

S. Sivabalan
C. Filsfils
Cisco
S. Litkowski
Orange
September 25, 2018

**Signaling Entropy Label Capability and Entropy Readable Label Depth
Using IS-IS
draft-ietf-isis-mpls-elc-06**

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 of processing 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), in the cases where stacked LSPs are used for whatever reasons. This document defines mechanisms to signal these two capabilities using IS-IS. These mechanisms are useful when the label advertisement is also done via IS-IS. In addition, this document introduces the Non-IGP Functional Capabilities Sub-TLV for advertising IS-IS router's actual non-IGP functional capabilities. ELC is one of such non-IGP functional capabilities.

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 29, 2019.

Copyright Notice

Copyright (c) 2018 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.	Non-IGP Functional Capabilities Sub-TLV	3
4.	Advertising ELC Using IS-IS	4
5.	Advertising ERLD Using IS-IS	4
6.	Acknowledgements	4
7.	IANA Considerations	5
8.	Security Considerations	5
9.	References	5
9.1.	Normative References	5
9.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

[RFC6790] describes a method to load balance Multiprotocol Label Switching (MPLS) traffic flows using Entropy Labels (EL). [RFC6790] introduces the concept of Entropy Label Capability (ELC) and defines the signalings of this capability via MPLS signaling protocols. Recently, mechanisms are being defined to signal labels via link-state Interior Gateway Protocols (IGP) such as IS-IS [I-D.ietf-isis-segment-routing-extensions]. In such scenario, the signaling mechanisms defined in [RFC6790] are inadequate. This draft defines a mechanism to signal the ELC [RFC6790] using IS-IS. This mechanism is useful when the label advertisement is also done via IS-IS.

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 whether it's necessary to insert an EL for a given LSP of the stacked 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](#)].

2. Terminology

This memo makes use of the terms defined in [[RFC6790](#)] and [[RFC4971](#)].

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](#)].

3. Non-IGP Functional Capabilities Sub-TLV

This document defines the Non-IGP Functional Capabilities Sub-TLV with Sub-TLV type of TBD1 within the body of the IS-IS Router Capability TLV. An IS-IS router advertising an IS-IS Router Capability TLV MAY include the Non-IGP Functional Capabilities Sub-TLV. The Sub-TLV MUST reflect the advertising IS-IS router's actual non-IGP functional capabilities in the flooding scope of the containing Router Capability TLV.

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

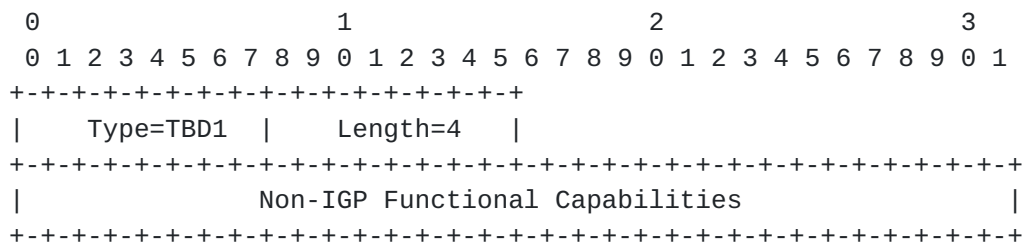


Figure 1: Non-IGP Functional Capabilities Sub-TLV Format

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-IGP 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-IGP Functional Capabilities Sub-TLV MAY be followed by optional Sub-TLVs that further specify a non-IGP functional capability. The specifications for non-IGP functional capabilities advertised in this Sub-TLV MUST describe protocol behavior and address backwards compatibility.

4. Advertising ELC Using IS-IS

One bit of the Non-IGP Functional Capability Bits (Bit 0 is desired) is to be assigned by the IANA for the ELC [[RFC6790](#)]. If a router has multiple line cards, the router MUST NOT announce the ELC [[RFC6790](#)] unless all of its linecards are capable of processing ELs.

How to apply the ELC advertisement to the inter-area, inter-AS and inter-protocol scenarios is outside the scope of this document.

5. Advertising ERLD Using IS-IS

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. As shown in Figure 2, it is formatted as described in [[I-D.ietf-isis-segment-routing-msd](#)] with a new MSD-Type code to be assigned by IANA (the type code of 2 is desired) and the Value field is set to the ERLD in the range between 0 to 255. The scope of the advertisement depends on the application. If a router has multiple linecards with different capabilities of reading the maximum label stack depth, the router MUST advertise the smallest one.

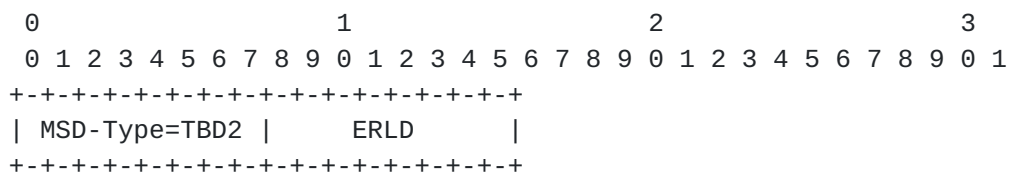


Figure 2: ERLD MSD-Type Format

6. 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.

7. IANA Considerations

This document requests IANA to allocate one sub-TLV type of the Router Capability TLV registry for the Non-IGP Functional Capabilities Sub-TLV. Furthermore, this document requests IANA to create a subregistry for "Non-IGP Functional Capability Bits" within the "Interior Gateway Protocol (IGP) Parameters" registry. This subregistry is comprised of the fields Bit Number, Capability Name, and Reference. Initially, one bit is requested to be assigned for the ELC. The registration procedure is "Expert Review" as defined in [RFC8126]. The following values are defined by this document:

Bit No.	Capability Name	Reference
-----	-----	-----
0	ELC	This document
1-31	Unassigned	This document

Figure 3: Non-IGP Functional Capability Bits Registry

IANA is requested to allocate a MSD type (the type code of 2 is desired) from the "IGP MSD Types" registry for ERLD.

8. Security Considerations

The security considerations as described in [RFC4971] is applicable to this document. This document does not introduce any new security risk.

9. References

9.1. Normative References

- [I-D.ietf-isis-segment-routing-extensions]
 Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-19](#) (work in progress), July 2018.
- [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-16](#) (work in progress), September 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-14](#) (work in progress), June 2018.

[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>>.

[RFC4971] Vasseur, JP., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), DOI 10.17487/RFC4971, July 2007, <<https://www.rfc-editor.org/info/rfc4971>>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

[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>>.

[RFC7794] Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability", [RFC 7794](#), DOI 10.17487/RFC7794, March 2016, <<https://www.rfc-editor.org/info/rfc7794>>.

[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>>.

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

Siva Sivabalan
Cisco

Email: msiva@cisco.com

Clarence Filsfils
Cisco

Email: cfilsfil@cisco.com

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com

