ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 9088
Category: Standards Track
ISSN: 2070-1721

X. Xu
Capitalonline
S. Kini

P. Psenak
C. Filsfils
S. Litkowski
Cisco Systems, Inc.
M. Bocci
Nokia

August 2021

Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS

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 Label Switched Path (LSP) unless an egress LSR has indicated via signaling that it has the capability to process ELs, referred to as the Entropy Label Capability (ELC), on that LSP. In addition, it would be useful for ingress LSRs to know each LSR's capability for 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 IS-IS and Border Gateway Protocol - Link State (BGP-LS).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9088.

Table of Contents

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 IS-IS [RFC8667]. This document defines a mechanism to signal the ELC using IS-IS.

In cases where Segment Routing (SR) is used with the MPLS data plane (e.g., SR-MPLS [RFC8660]), 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 [RFC8662], 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. This document defines a mechanism to signal the ERLD using IS-IS.

2. Terminology

This memo makes use of the terms defined in [RFC6790], and [RFC8662].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Advertising ELC Using IS-IS

Even though ELC is a property of the node, in some cases it is advantageous to associate and advertise the ELC with a prefix. In a multi-area network, 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.

Bit 3 in the Prefix Attribute Flags [RFC7794] is used as the ELC Flag (E-Flag), as shown in Figure 1. If a router has multiple interfaces, the router MUST NOT announce the ELC for any local host prefixes unless all of its interfaces are capable of processing ELs. If a router supports ELs on all of its interfaces, it SHOULD set the ELC for every local host prefix it advertises in IS-IS.

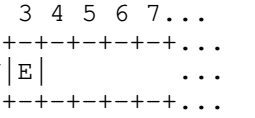

Figure 1: Prefix Attribute Flags

E-Flag:
ELC Flag (Bit 3) - Set for local host prefix of the originating node if it supports ELC on all interfaces.

The ELC signaling MUST be preserved when a router propagates a prefix between IS-IS levels [RFC5302].

When redistributing a prefix between two IS-IS protocol instances or redistributing from another protocol to an IS-IS protocol instance, a

router SHOULD preserve the ELC signaling for that prefix if it exists. The exact mechanism used to exchange ELC between protocol instances running on an Autonomous System Border Router is outside of the scope of this document.

4. Advertising ERLD Using IS-IS

A new MSD-Type [RFC8491], called ERLD-MSD, is defined to advertise the ERLD [RFC8662] of a given router. An MSD-Type code 2 has been assigned by IANA for ERLD-MSD. The MSD-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 interfaces with different capabilities of reading the maximum label stack depth, the router MUST advertise the smallest value found across all its interfaces.

The absence of ERLD-MSD advertisements indicates only that the advertising node does not support advertisement of this capability.

The considerations for advertising the ERLD are specified in [RFC8662].

If the ERLD-MSD type is received in the Link MSD sub-TLV, it MUST be ignored.

5. Signaling ELC and ERLD in BGP-LS

The IS-IS extensions defined in this document can be advertised via BGP-LS (distribution of Link-State and TE information using BGP) [RFC7752] using existing BGP-LS TLVs.

The ELC is advertised using the Prefix Attribute Flags TLV as defined in [RFC9085].

The ERLD-MSD is advertised using the Node MSD TLV as defined in [RFC8814].

6. IANA Considerations

IANA has completed the following actions for this document:

* Bit 3 in the "Bit Values for Prefix Attribute Flags Sub-TLV" registry has been assigned to the ELC Flag. IANA has updated the registry to reflect the name used in this document: ELC Flag (E-Flag).

* Type 2 in the "IGP MSD-Types" registry has been assigned for the ERLD-MSD. IANA has updated the registry to reflect the name used in this document: ERLD-MSD.

7. Security Considerations

This document specifies the ability to advertise additional node capabilities using IS-IS and BGP-LS. As such, the security considerations as described in [RFC7752], [RFC7794], [RFC7981], [RFC8491], [RFC8662], [RFC8814], and [RFC9085] are applicable to this document.

Incorrectly setting the E-Flag during origination, propagation, or redistribution may lead to poor or no load-balancing of the MPLS traffic or to MPLS traffic being discarded on the egress node.

Incorrectly setting the ERLD value may lead to poor or no load-balancing of the MPLS traffic.

8. References

8.1. 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,
<https://www.rfc-editor.org/info/rfc2119>.

[RFC5302] Li, T., Smit, H., and T. Przygienda, "Domain-Wide Prefix Distribution with Two-Level IS-IS", RFC 5302, DOI 10.17487/RFC5302, October 2008, <https://www.rfc-editor.org/info/rfc5302>.

[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>.

[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>.

[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>.

[RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <https://www.rfc-editor.org/info/rfc7981>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <https://www.rfc-editor.org/info/rfc8491>.

[RFC8662] Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and J. Tantsura, "Entropy Label for Source Packet Routing in Networking (SPRING) Tunnels", RFC 8662, DOI 10.17487/RFC8662, December 2019, <https://www.rfc-editor.org/info/rfc8662>.

[RFC8814] Tantsura, J., Chunduri, U., Talaulikar, K., Mirsky, G., and N. Triantafillis, "Signaling Maximum SID Depth (MSD) Using the Border Gateway Protocol - Link State", RFC 8814, DOI 10.17487/RFC8814, August 2020, <https://www.rfc-editor.org/info/rfc8814>.

[RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <https://www.rfc-editor.org/info/rfc9085>.

8.2. Informative References

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <https://www.rfc-editor.org/info/rfc8660>.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <https://www.rfc-editor.org/info/rfc8667>.

Acknowledgements

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

Contributors

The following people contributed to the content of this document and should be considered as coauthors:

Gunter Van de Velde (editor)
Nokia
Antwerp
Belgium

Email: gunter.van_de_velde@nokia.com

Wim Henderickx
Nokia
Belgium

Email: wim.henderickx@nokia.com

Keyur Patel
Arrcus
United States of America

Email: keyur@arrcus.com

Authors' Addresses

Xiaohu Xu
Capitalonline

Email: xiaohu.xu@capitalonline.net

Sriganesh Kini

Email: sriganeshkini@gmail.com

Peter Psenak
Cisco Systems, Inc.
Eurovea Centre, Central 3
Pribinova Street 10
81109 Bratislava
Slovakia

Email: ppsenak@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

Stephane Litkowski
Cisco Systems, Inc.
La Rigourdiere
Cesson Sevigne
France

Email: slitkows@cisco.com

Matthew Bocci
Nokia
740 Waterside Drive
Aztec West Business Park
Bristol
BS32 4UF
United Kingdom

Email: matthew.bocci@nokia.com