IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 23, 2019

J. Tantsura
Apstra, Inc.
U. Chunduri
Huawei USA
G. Mirsky
ZTE Corp.
S. Sivabalan
Cisco
N. Triantafillis
Apstra, Inc.
February 19, 2019

Signaling MSD (Maximum SID Depth) using Border Gateway Protocol Link-State
draft-ietf-idr-bgp-ls-segment-routing-msd-04

Abstract

This document defines a way for a Border Gateway Protocol Link-State (BGP-LS) speaker to advertise multiple types of supported Maximum SID Depths (MSDs) at node and/or link granularity.

Such advertisements allow logically centralized entities (e.g., centralized controllers) to determine whether a particular SID stack can be supported in a given network.

Table of Contents

1. Introduction

When Segment Routing tunnels are computed by a centralized controller, it is critical that the controller learns the MSD "Maximum SID Depth" of the node or link SR tunnel exits over, so the SID stack depth of a path computed doesn't exceed the number of SIDs the node is capable of imposing. This document describes how to use BGP-LS to signal the MSD of a node or link to a centralized controller.

PCEP SR extensions draft [I-D.ietf-pce-segment-routing] signals MSD in SR PCE Capability TLV and METRIC Object. However, if PCEP is not supported/configured on the head-end of a SR tunnel or a Binding-SID anchor node and controller does not participate in IGP routing, it has no way to learn the MSD of nodes and links which has been configured. BGP-LS [RFC7752] defines a way to expose topology and

associated attributes and capabilities of the nodes in that topology to a centralized controller.

Other types of MSD are known to be useful. For example, [I-D.ietf-ospf-mpls-elc] and [I-D.ietf-isis-mpls-elc] define Readable Label Depth Capability (RLDC) that is used by a head-end to insert an Entropy Label (EL) at a depth that can be read by transit nodes.

1.1. Conventions used in this document

1.1.1. Terminology

BGP-LS: Distribution of Link-State and TE Information using Border Gateway Protocol

MSD: Maximum SID Depth

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol

SID: Segment Identifier

SR: Segment routing

1.1.2. Requirements Language

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 .

2. Problem Statement

In existing technology only PCEP has extension to signal the MSD (SR PCE Capability TLV/ METRIC Object as defined in [I-D.ietf-pce-segment-routing],If PCEP is not supported by the node (head-end of the SR tunnel) controller has no way to learn the MSD of the node/link configured. OSPF and IS-IS extensions are defined in:

[RFC8476], [RFC8491]

3. MSD supported by a node

Node MSD is encoded in a new Node Attribute TLV, as defined in [RFC7752]

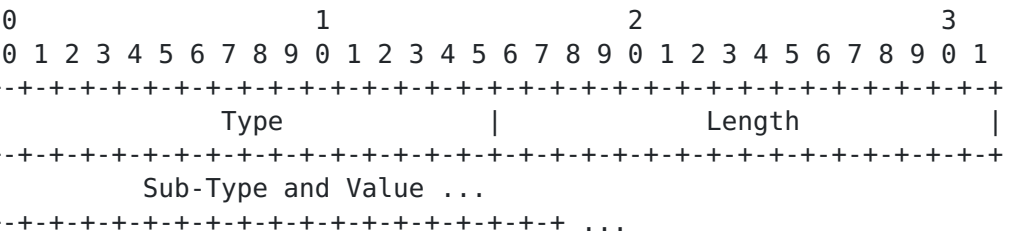

Figure 1: Node attribute format

Type : A 2-octet field specifying code-point of the new TLV type. Code-point:(TBD1) from BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs registry

Length: A 2-octet field that indicates the length of the value portion

Sub-Type and value fields are as defined in corresponding OSPF [RFC8476] and IS-IS [RFC8491] extensions.

4. MSD supported on a link

Link MSD is encoded in a New Link Attribute TLV, as defined in [RFC7752]

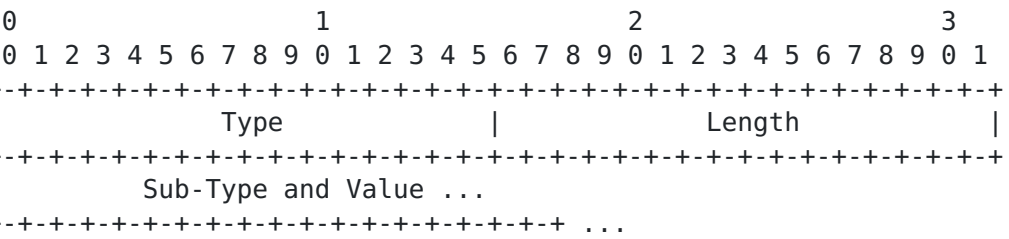

Figure 2: Link attribute format

Type : A 2-octet field specifying code-point of the new TLV type. Code-point:(TBD2) from BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs registry

Length: A 2-octet field that indicates the length of the value portion

Sub-Type and value fields are as defined in corresponding OSPF [RFC8476] and IS-IS [RFC8491] extensions.

5. IANA Considerations

We request IANA assign code points from the registry BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs, as follows: TLV Code Point Description IS-IS TLV/Sub-TLV Reference TBD1 Node MSD 242/23 (this document) TBD2 Link MSD (22,23,25,141,222,223)/15 (this document)

6. Security Considerations

The advertisement of an incorrect MSD value may have negative consequences. If the value is smaller than supported, path computation may fail to compute a viable path. If the value is larger than supported, an attempt to instantiate a path that can't be supported by the head-end (the node performing the SID imposition) may occur. The presence of this information may also inform an attacker of how to induce any of the aforementioned conditions.

This document does not introduce security issues beyond those discussed in [RFC7752], [RFC8476] and [RFC8491]

7. Acknowledgements

We like to thank Acee Lindem, Ketan Talaulikar, Stephane Litkowski and Bruno Decraene for their reviews and valuable comments.

8. References

8.1. Normative References

[I-D.ietf-pce-segment-routing]
Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", draft-ietf-pce-segment-routing-15 (work in progress), February 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>.

[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>.

[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>.

[RFC8476] Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak, "Signaling Maximum SID Depth (MSD) Using OSPF", RFC 8476, DOI 10.17487/RFC8476, December 2018, <https://www.rfc-editor.org/info/rfc8476>.

[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>.

8.2. Informative References

[I-D.ietf-isis-mpls-elc]
Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS", draft-ietf-isis-mpls-elc-06 (work in progress), September 2018.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-22 (work in progress), December 2018.

[I-D.ietf-ospf-mpls-elc]
Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Entropy Readable Label-stack Depth Using OSPF", draft-ietf-ospf-mpls-elc-07 (work in progress), September 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-18 (work in progress), December 2018.

Authors' Addresses

Jeff Tantsura
Apstra, Inc.

Email: jefftant.ietf@gmail.com

Uma Chunduri
Huawei USA

Email: uma.chunduri@huawei.com

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

Siva Sivabalan
Cisco

Email: msiva@cisco.com

Nikos Triantafillis
Apstra, Inc.

Email: nikos@apstra.com