

Internet Engineering Task Force (IETF)
Request for Comments: 8491
Category: Standards Track
ISSN: 2070-1721

J. Tantsura
Apstra, Inc.
U. Chunduri
Huawei Technologies
S. Aldrin
Google, Inc.
L. Ginsberg
Cisco Systems
November 2018

Signaling Maximum SID Depth (MSD) Using IS-IS

Abstract

This document defines a way for an Intermediate System to Intermediate System (IS-IS) router to advertise multiple types of supported Maximum SID Depths (MSDs) at node and/or link granularity. Such advertisements allow entities (e.g., centralized controllers) to determine whether a particular Segment ID (SID) stack can be supported in a given network. This document only defines one type of MSD: Base MPLS Imposition. However, it defines an encoding that can support other MSD types. This document focuses on MSD use in a network that is Segment Routing (SR) enabled, but MSD may also be useful when SR is not enabled.

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/rfc8491>.

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
1.1. Terminology	3
1.2. Requirements Language	4
2. Node MSD Advertisement	4
3. Link MSD Advertisement	5
4. Procedures for Defining and Using Node and Link MSD Advertisements	6
5. Base MPLS Imposition MSD	6
6. IANA Considerations	7
7. Security Considerations	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Acknowledgements	10
Contributors	10
Authors' Addresses	10

1. Introduction

When Segment Routing (SR) paths are computed by a centralized controller, it is critical that the controller learn the Maximum SID Depth (MSD) that can be imposed at each node/link of a given SR path. This ensures that the Segment Identifier (SID) stack depth of a computed path does not exceed the number of SIDs the node is capable of imposing.

[PCEP-EXT] defines how to signal MSD in the Path Computation Element Communication Protocol (PCEP). However, if PCEP is not supported/configured on the head-end of an SR tunnel or a Binding-SID anchor node, and the controller does not participate in IGP routing, it has no way of learning the MSD of nodes and links. BGP-LS (Distribution

of Link-State and TE Information Using Border Gateway Protocol) [RFC7752] defines a way to expose topology and associated attributes and capabilities of the nodes in that topology to a centralized controller. MSD signaling by BGP-LS has been defined in [MSD-BGP]. Typically, BGP-LS is configured on a small number of nodes that do not necessarily act as head-ends. In order for BGP-LS to signal MSD for all the nodes and links in the network for which MSD is relevant, MSD capabilities SHOULD be advertised by every Intermediate System to Intermediate System (IS-IS) router in the network.

Other types of MSDs are known to be useful. For example, [ELC-ISIS] defines Entropy Readable Label Depth (ERLD), which is used by a head-end to insert an Entropy Label (EL) at a depth where it can be read by transit nodes.

This document defines an extension to IS-IS used to advertise one or more types of MSDs at node and/or link granularity. It also creates an IANA registry for assigning MSD-Type identifiers and defines the Base MPLS Imposition MSD-Type. In the future, it is expected that new MSD-Types will be defined to signal additional capabilities, e.g., entropy labels, SIDs that can be imposed through recirculation, or SIDs associated with another data plane such as IPv6.

MSD advertisements MAY be useful even if Segment Routing itself is not enabled. For example, in a non-SR MPLS network, MSD defines the maximum label depth.

1.1. Terminology

BMI: Base MPLS Imposition is the number of MPLS labels that can be imposed inclusive of all service/transport/special labels.

MSD: Maximum SID Depth is the number of SIDs supported by a node or a link on a node.

SID: Segment Identifier is defined in [RFC8402].

Label Imposition: Imposition is the act of modifying and/or adding labels to the outgoing label stack associated with a packet. This includes:

- * replacing the label at the top of the label stack with a new label
- * pushing one or more new labels onto the label stack

The number of labels imposed is then the sum of the number of labels that are replaced and the number of labels that are pushed. See [RFC3031] for further details.

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. Node MSD Advertisement

The Node MSD sub-TLV is defined within the body of the IS-IS Router CAPABILITY TLV [RFC7981] to carry the provisioned SID depth of the router originating the IS-IS Router CAPABILITY TLV. Node MSD is the smallest MSD supported by the node on the set of interfaces configured for use by the advertising IGP instance. MSD values may be learned via a hardware API or may be provisioned.




Figure 1: Node MSD Sub-TLV

Type: 23

Length: variable (multiple of 2 octets); represents the total length of the Value field

Value: field consists of one or more pairs of a 1-octet MSD-Type and 1-octet MSD-Value

MSD-Type: value defined in the "IGP MSD-Types" registry created by the IANA Considerations section of this document Section 6

MSD-Value: number in the range of 0-255 (for all MSD-Types, 0 represents the lack of ability to support a SID stack of any depth; any other value represents that of the node. This value MUST represent the lowest value supported by any link configured for use by the advertising IS-IS instance.)

This sub-TLV is optional. The scope of the advertisement is specific to the deployment.

If there exist multiple Node MSD advertisements for the same MSD-Type originated by the same router, the procedures defined in [RFC7981] apply. These procedures may result in different MSD values being used, for example, by different controllers. This does not, however, create any interoperability issue.

3. Link MSD Advertisement

The Link MSD sub-TLV is defined for TLVs 22, 23, 25, 141, 222, and 223 to carry the MSD of the interface associated with the link. MSD values may be signaled by the forwarding plane or may be provisioned.




Figure 2: Link MSD Sub-TLV

Type: 15

Length: variable (multiple of 2 octets); represents the total length of the Value field

Value: field consists of one or more pairs of a 1-octet MSD-Type and 1-octet MSD-Value

MSD-Type: value defined in the "IGP MSD-Types" registry created by the IANA Considerations section of this document Section 6

MSD-Value: number in the range of 0-255 (for all MSD-Types, 0 represents the lack of ability to support a SID stack of any depth; any other value represents that of the particular link when used as an outgoing interface.)

This sub-TLV is optional.

If multiple Link MSD advertisements for the same MSD-Type and the same link are received, the procedure to select which copy to use is undefined.

If the advertising router performs label imposition in the context of the ingress interface, it is not possible to meaningfully advertise per-link values. In such a case, only the Node MSD SHOULD be advertised.

4. Procedures for Defining and Using Node and Link MSD Advertisements

When Link MSD is present for a given MSD-Type, the value of the Link MSD MUST take precedence over the Node MSD. If a Link MSD-Type is not signaled, but the Node MSD-Type is, then the Node MSD-Type value MUST be considered to be the MSD value for that link.

In order to increase flooding efficiency, it is RECOMMENDED that routers with homogenous Link MSD values advertise just the Node MSD value.

The meaning of the absence of both Node and Link MSD advertisements for a given MSD-Type is specific to the MSD-Type. Generally, it can only be inferred that the advertising node does not support advertisement of that MSD-Type. In some cases, however, the lack of advertisement might imply that the functionality associated with the MSD-Type is not supported. The correct interpretation MUST be specified when an MSD-Type is defined.

5. Base MPLS Imposition MSD

Base MPLS Imposition MSD (BMI-MSD) signals the total number of MPLS labels that can be imposed, including all service/transport/special labels.

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

6. IANA Considerations

IANA has allocated a sub-TLV type for the new sub-TLV proposed in Section 2 of this document from the "Sub-TLVs for TLV 242 (IS-IS Router CAPABILITY TLV)" registry as defined by [RFC7981].

IANA has allocated the following value:

Value	Description	Reference
-----	-----	-----
23	Node MSD	This document

Figure 3: Node MSD

IANA has allocated a sub-TLV type as defined in Section 3 from the "Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 (Extended IS reachability, IS Neighbor Attribute, L2 Bundle Member Attributes, inter-AS reachability information, MT-ISN, and MT IS Neighbor Attribute TLVs)" registry.

IANA has allocated the following value:

Value	Description	Reference
-----	-----	-----
15	Link MSD	This document

Figure 4: Link MSD

Per-TLV information where Link MSD sub-TLV can be part of:

TLV	22	23	25	141	222	223
---	-----	-----	-----	-----	-----	-----
	Y	Y	Y	Y	Y	Y

Figure 5: TLVs Where LINK MSD Sub-TLV Can Be Present

IANA has created an IANA-managed registry titled "IGP MSD-Types" under the "Interior Gateway Protocol (IGP) Parameters" registry to identify MSD-Types as proposed in Sections 2 and 3. The registration procedure is "Expert Review" as defined in [RFC8126]. Types are an unsigned 8-bit number. The following values are defined by this document:

Value	Name	Reference
-----	-----	-----
0	Reserved	This document
1	Base MPLS Imposition MSD	This document
2-250	Unassigned	
251-254	Experimental Use	This document
255	Reserved	This document

Figure 6: MSD-Types Codepoints Registry

General guidance for the designated experts is defined in [RFC7370].

7. Security Considerations

Security considerations as specified by [RFC7981] are applicable to this document.

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.

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>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.

- [RFC7370] Ginsberg, L., "Updates to the IS-IS TLV Codepoints Registry", RFC 7370, DOI 10.17487/RFC7370, September 2014, <<https://www.rfc-editor.org/info/rfc7370>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

8.2. Informative References

- [ELC-ISIS] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS", Work in Progress, draft-ietf-isis-mpls-elc-06, September 2018.
- [MSD-BGP] Tantsura, J., Chunduri, U., Mirsky, G., and S. Sivabalan, "Signaling MSD (Maximum SID Depth) using Border Gateway Protocol Link-State", Work in Progress, draft-ietf-idr-bgp-ls-segment-routing-msd-02, August 2018.
- [PCEP-EXT] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", Work in Progress, draft-ietf-pce-segment-routing-13, October 2018.
- [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>>.

Acknowledgements

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

Contributors

The following people contributed to this document:

Peter Psenak

Email: ppsenak@cisco.com

Authors' Addresses

Jeff Tantsura
Apstra, Inc.

Email: jefftant.ietf@gmail.com

Uma Chunduri
Huawei Technologies

Email: uma.chunduri@huawei.com

Sam Aldrin
Google, Inc.

Email: aldrin.ietf@gmail.com

Les Ginsberg
Cisco Systems

Email: ginsberg@cisco.com

