

Internet Engineering Task Force (IETF)
Request for Comments: 8476
Category: Standards Track
ISSN: 2070-1721

J. Tantsura
Apstra, Inc.
U. Chunduri
Huawei Technologies
S. Aldrin
Google, Inc.
P. Psenak
Cisco Systems
December 2018

Signaling Maximum SID Depth (MSD) Using OSPF

Abstract

This document defines a way for an Open Shortest Path First (OSPF) 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 Identifier (SID) stack can be supported in a given network. This document only refers to the Signaling MSD as defined in RFC 8491, but it defines an encoding that can support other MSD types. Here, the term "OSPF" means both OSPFv2 and OSPFv3.

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/rfc8476>.

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	3
1.1. Terminology	4
1.2. Requirements Language	4
2. Node MSD Advertisement	5
3. Link MSD Sub-TLV	6
4. Procedures for Defining and Using Node and Link MSD Advertisements	7
5. IANA Considerations	7
6. Security Considerations	8
7. References	9
7.1. Normative References	9
7.2. Informative References	10
Acknowledgements	11
Contributors	11
Authors' Addresses	11

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 on a given SR path. This ensures that the Segment Identifier (SID) stack depth of a computed path doesn't 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 BGP) [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 OSPF 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 OSPF used to advertise one or more types of MSDs at node and/or link granularity. In the future, it is expected that new MSD-Types will be defined to signal additional capabilities, e.g., ELs, SIDs that can be imposed through recirculation, or SIDs associated with another data plane such as IPv6.

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

1.1. Terminology

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

BGP-LS: Distribution of Link-State and TE Information Using BGP

OSPF: Open Shortest Path First

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

SID: Segment Identifier as 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.

PCEP: Path Computation Element Communication Protocol

SR: Segment Routing

LSA: Link State Advertisement

RI: Router Information

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 TLV within the body of the OSPF RI Opaque LSA [RFC7770] is defined to carry the provisioned SID depth of the router originating the RI LSA. 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 TLV

Type: 12

Length: variable (multiple of 2 octets); represents the total length of the value field in octets.

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

MSD-Type: one of the values defined in the "IGP MSD-Types" registry defined in [RFC8491].

MSD-Value: a number in the range of 0-255. For all MSD-Types, 0 represents the lack of ability to impose an MSD 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 OSPF instance.

This TLV is optional and is applicable to both OSPFv2 and OSPFv3. The scope of the advertisement is specific to the deployment.

When multiple Node MSD TLVs are received from a given router, the receiver MUST use the first occurrence of the TLV in the Router Information (RI) LSA. If the Node MSD TLV appears in multiple RI LSAs that have different flooding scopes, the Node MSD TLV in the RI LSA with the area-scoped flooding scope MUST be used. If the Node MSD TLV appears in multiple RI LSAs that have the same flooding scope, the Node MSD TLV in the RI LSA with the numerically smallest Instance ID MUST be used and other instances of the Node MSD TLV MUST be ignored. The RI LSA can be advertised at any of the defined

opaque flooding scopes (link, area, or Autonomous System (AS)). For the purpose of Node MSD TLV advertisement, area-scoped flooding is RECOMMENDED.

3. Link MSD Sub-TLV

The Link MSD sub-TLV is defined to carry the MSD of the interface associated with the link. MSD values may be learned via a hardware API or may be provisioned.




Figure 2: Link MSD Sub-TLV

Type:

For OSPFv2, the link-level MSD-Value is advertised as an optional sub-TLV of the OSPFv2 Extended Link TLV as defined in [RFC7684] and has a type of 6.

For OSPFv3, the link-level MSD-Value is advertised as an optional sub-TLV of the E-Router-LSA TLV as defined in [RFC8362] and has a type of 9.

Length: variable; same as defined in Section 2.

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

MSD-Type: one of the values defined in the "IGP MSD-Types" registry defined in [RFC8491].

The MSD-Value field contains the Link MSD of the router originating the corresponding LSA as specified for OSPFv2 and OSPFv3. The Link MSD is a number in the range of 0-255. For all MSD-Types, 0 represents the lack of ability to impose an MSD stack of any depth; any other value represents that of the particular link when used as an outgoing interface.

If this sub-TLV is advertised multiple times for the same link in different OSPF Extended Link Opaque LSAs / E-Router-LSAs originated by the same OSPF router, the sub-TLV in the OSPFv2 Extended Link

Opaque LSA with the smallest Opaque ID or in the OSPFv3 E-Router-LSA with the smallest Link State ID MUST be used by receiving OSPF routers. This situation SHOULD be logged as an error.

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. When a Link MSD-Type is not signaled but the Node MSD-Type is, then the Node MSD-Type value MUST be considered as 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. However, in some cases the lack of advertisement might imply that the functionality associated with the MSD-Type is not supported. Per [RFC8491], the correct interpretation MUST be specified when an MSD-Type is defined.

5. IANA Considerations

This specification updates several existing OSPF registries.

IANA has allocated TLV type 12 from the "OSPF Router Information (RI) TLVs" registry as defined by [RFC7770].

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

Figure 3: RI Node MSD

IANA has allocated sub-TLV type 6 from the "OSPFv2 Extended Link TLV Sub-TLVs" registry.

Value	Description	Reference
-----	-----	-----
6	OSPFv2 Link MSD	This document

Figure 4: OSPFv2 Link MSD

IANA has allocated sub-TLV type 9 from the "OSPFv3 Extended-LSA Sub-TLVs" registry.

Value	Description	Reference
-----	-----	-----
9	OSPFv3 Link MSD	This document

Figure 5: OSPFv3 Link MSD

6. Security Considerations

Security concerns for OSPF are addressed in [RFC7474], [RFC4552], and [RFC7166]. Further security analysis for the OSPF protocol is done in [RFC6863]. Security considerations as specified by [RFC7770], [RFC7684], and [RFC8362] are applicable to this document.

Implementations MUST ensure that malformed TLVs and sub-TLVs defined in this document are detected and do not provide a vulnerability for attackers to crash the OSPF router or routing process. Reception of malformed TLVs or sub-TLVs SHOULD be counted and/or logged for further analysis. Logging of malformed TLVs and sub-TLVs SHOULD be rate-limited to prevent a Denial-of-Service (DoS) attack (distributed or otherwise) from overloading the OSPF control plane.

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.

There's no DoS risk specific to this extension, and it is not vulnerable to replay attacks.

7. References

7.1. 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>>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.
- [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, <<https://www.rfc-editor.org/info/rfc7770>>.
- [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>>.
- [RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.
- [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>>.
- [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>>.

7.2. Informative References

- [ELC-ISIS] Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Entropy Readable Label-stack Depth Using OSPF", Work in Progress, draft-ietf-ospf-mpls-elc-07, 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-14, October 2018.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [RFC6863] Hartman, S. and D. Zhang, "Analysis of OSPF Security According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6863, DOI 10.17487/RFC6863, March 2013, <<https://www.rfc-editor.org/info/rfc6863>>.
- [RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 7166, DOI 10.17487/RFC7166, March 2014, <<https://www.rfc-editor.org/info/rfc7166>>.
- [RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <<https://www.rfc-editor.org/info/rfc7474>>.
- [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, Tal Mizrahi, Stephane Litkowski, and Bruno Decraene for their reviews and valuable comments.

Contributors

The following person contributed to this document:

Les Ginsberg

Email: ginsberg@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

Peter Psenak
Cisco Systems

Email: ppsenak@cisco.com

