

OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 4, 2019

J. Tantsura
Nuage Networks
U. Chunduri
Huawei Technologies
S. Aldrin
Google, Inc
P. Psenak
Cisco Systems
August 31, 2018

Signaling MSD (Maximum SID Depth) using OSPF
draft-ietf-ospf-segment-routing-msd-20

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 SID stack can be supported in a given network. This document defines only one type of MSD, but defines an encoding that can support other MSD types. Here the term OSPF means both OSPFv2 and OSPFv3.

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 4, 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
1.1.	Terminology	3
1.2.	Requirements Language	4
2.	Node MSD Advertisement	4
3.	Link MSD sub-TLV	5
4.	Using Node and Link MSD Advertisements	6
5.	IANA Considerations	7
6.	Security Considerations	7
7.	Contributors	7
8.	Acknowledgments	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	8
	Authors' Addresses	10

[1.](#) Introduction

When Segment Routing (SR) paths are computed by a centralized controller, it is critical that the controller learns the Maximum SID Depth (MSD) that can be imposed at each node/link on a given SR path to insure that the SID stack depth of a computed path doesn't exceed the number of SIDs the node is capable of imposing.

Path Computation Element Protocol(PCEP) SR draft [[I-D.ietf-pce-segment-routing](#)] signals MSD in SR Path Computation Element Capability TLV and METRIC Object. However, if PCEP is not supported/configured on the head-end of an 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. 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 [[I-D.ietf-idr-bgp-ls-segment-routing-msd](#)]. 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 where MSD is relevant, MSD capabilities should be advertised by every OSPF router in the network.

Other types of MSD are known to be useful. For example, [\[I-D.ietf-ospf-mpls-elc\]](#) defines 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.

This document defines an extension to OSPF used to advertise one or more types of MSD 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., entropy labels, SIDs that can be imposed through recirculation, or SIDs associated with another dataplane e.g., IPv6. Although MSD advertisements are associated with Segment Routing, the advertisements MAY be present even if Segment Routing itself is not enabled. Note that in a non-SR MPLS network, label depth is what is defined by the MSD advertisements.

1.1. Terminology

This memo makes use of the terms defined in [\[RFC7770\]](#)

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

OSPF: Open Shortest Path First

MSD: Maximum SID Depth - the number of SIDs a node or one of its links can support

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol

SR: Segment Routing

SID: Segment Identifier

LSA: Link state advertisement

RI: OSPF Router Information LSA

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.

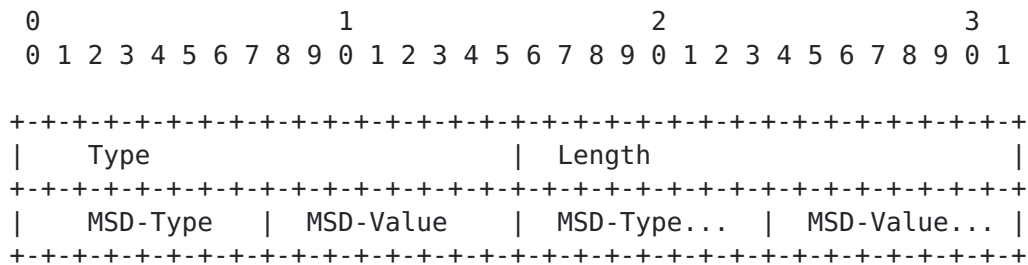


Figure 1: Node MSD TLV

Type: TBD1

Length: variable (multiple of 2 octets) and represents the total length of 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 [\[I-D.ietf-isis-segment-routing-msd\]](#).

MSD-Value: a number in the range of 0-255. For all MSD-Types, 0 represents lack of the ability to impose 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 applicable to OSPFv2 and to OSPFv3 and is optional. 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 LSA. If the Node MSD TLV appears in multiple Router Information LSAs that have different flooding scopes, the Node MSD TLV in the Router Information LSA with the area-scoped flooding scope MUST be used. If the Node MSD TLV appears in multiple Router Information LSAs that have the same flooding scope, the Node MSD TLV in the Router Information (RI) LSA with the numerically smallest Instance ID MUST be used and subsequent 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 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.

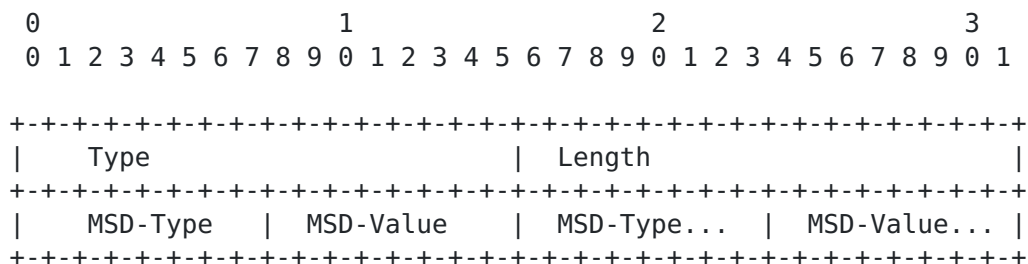


Figure 2: Link MSD Sub-TLV

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 TBD2.

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 TBD3.

Length: variable and 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 MSD Types registry defined in [[I-D.ietf-isis-segment-routing-msd](#)].

MSD-Value field contains Link MSD of the router originating the corresponding LSA as specified for OSPFv2 and OSPFv3. Link MSD is a number in the range of 0-255. For all MSD-Types, 0 represents lack of the ability to impose 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 in the same OSPFv2 Extended Link Opaque LSA/E-Router-LSA, only the first instance of the TLV MUST be used by receiving OSPF routers. This situation SHOULD be logged as an error.

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 OSPFv2 Extended Link 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 MAY be logged as a warning.

4. Using Node and Link MSD Advertisements

When Link MSD is present for a given MSD type, the value of the Link MSD MUST take preference over the Node MSD. When a Link MSD type is not signalled but the Node MSD type is, then the value of that Node MSD type MUST be considered as the corresponding Link MSD type value. In order to increase flooding efficiency, it is RECOMMENDED, that routers with homogenous Link MSD values advertise just the Node MSD value.

Information received in an MSD advertisements is to ensure that the controller learns the Maximum SID Depth (MSD) that can be imposed at each node/link on a given SR path so that the SID stack depth of a computed path doesn't exceed the number of SIDs the node is capable of imposing

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. The correct interpretation MUST be specified when an MSD type is defined.

5. IANA Considerations

This document requests IANA to allocate TLV type (TBD1) from the OSPF Router Information (RI) TLVs Registry as defined by [\[RFC7770\]](#). IANA has allocated the value 12 through the early assignment process. Also, this document requests IANA to allocate a sub-TLV type (TBD2) from the OSPFv2 Extended Link TLV Sub-TLVs registry. IANA has allocated the value 6 through the early assignment process. Finally, this document requests IANA to allocate a sub-TLV type (TBD3) from the OSPFv3 Extended-LSA Sub-TLV registry.

6. Security Considerations

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

Implementations MUST assure that malformed TLV and Sub-TLV 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 TLV or Sub-TLV 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 result:

If the value is smaller than supported - path computation failing to compute a viable path.

If the value is larger than supported - instantiation of a path that can't be supported by the head-end (the node performing the SID imposition).

The MSD discloses capabilities of the nodes (how many SIDs it supports), which could provide an indication of the abilities or even types of the nodes being used. This information could be used to gain intelligence about devices in the network.

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

7. Contributors

The following people contributed to this document:

Les Ginsberg

Email: ginsberg@cisco.com

8. Acknowledgments

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

9. References

9.1. Normative References

- [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-13](#) (work in progress), July 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>>.
- [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>>.

9.2. Informative References

- [I-D.ietf-idr-bgp-ls-segment-routing-msd]
Tantsura, J., Chunduri, U., Mirsky, G., and S. Sivabalan,
"Signaling MSD (Maximum SID Depth) using Border Gateway
Protocol Link-State", [draft-ietf-idr-bgp-ls-segment-
routing-msd-02](#) (work in progress), August 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-06](#) (work in progress), August 2018.
- [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-12](#) (work in progress), June
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>>.

Authors' Addresses

Jeff Tantsura
Nuage Networks

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