OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 7, 2018

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

April 05, 2018

Signaling MSD (Maximum SID Depth) using OSPF
draft-ietf-ospf-segment-routing-msd-10

Abstract

This document defines a way for an 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 only defines one type of MSD maximum label imposition, but defines an encoding which 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 October 7, 2018.

Copyright Notice

Table of Contents

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 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.

The 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 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 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. 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 MSD is relevant, MSD capabilites should be advertised to 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. It also creates an IANA registry for assigning MSD type identifiers. It laso 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 dataplane e.g., IPv6.

1.1. Conventions used in this document

1.1.1. Terminology

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

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

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

SID: Segment Identifier

SR: Segment Routing

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

BCP14 [RFC2119], [RFC8174] when, and only when they appear in all capitals, as shown here .

2. Terminology

This memo makes use of the terms defined in [RFC4970].

3. Node MSD TLV

The node MSD TLV within the body of the OSPF RI Opaque LSA is defined to carry the provisioned SID depth of the router originating the RI LSA. Node MSD is the minimum MSD supported by the node.

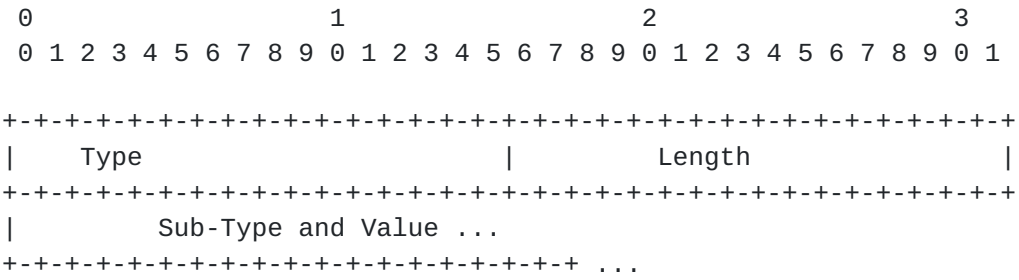

Figure 1: Node MSD TLV

The Type: TBD1

Length: variable (minimum of 2, multiple of 2 octets) and represents the total length of value field.

Value: consists of a 1 octet sub-type (IANA Registry) and 1 octet value.

Sub-Type 1 (IANA Section), MSD and the Value field contains maximum MSD of the router originating the RI LSA. Node Maximum MSD is a number in the range of 0-254. 0 represents lack of the ability to impose MSD stack of any depth; any other value represents that of the node. This value SHOULD represent the minimum value supported by a node.

Other Sub-types other than defined above are reserved for future extensions.

This TLV is applicable to OSPFv2 and to OSPFv3 [RFC5838] and is optional. The scope of the advertisement is specific to the deployment.

4. Link MSD sub-TLV

A new sub-TLV called Link MSD sub-TLV is defined to carry the provisioned SID depth of the interface associated with the link.

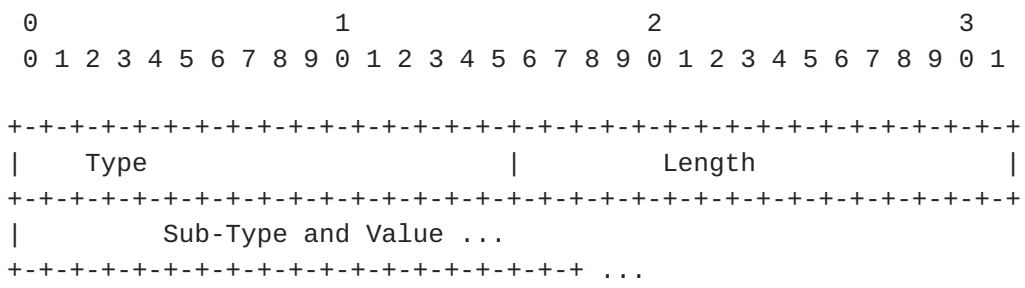

Figure 2: Link MSD Sub-TLV

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 value of TBD2.

For OSPFv3, the Link level MSD value is advertised as an optional Sub-TLV of the Router-Link TLV as defined in [I-D.ietf-ospf-ospfv3-lsa-extend], and has value of TBD3.

Length: variable and similar to what is defined in Section 3.

Value: consists of a 1 octet sub-type (IANA Registry) and 1 octet value.

Sub-Type 1 (IANA Section), MSD and the 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 0-254. 0 represents lack of the ability to impose MSD stack of any depth; any other value represents that of the particular link MSD value.

Other Sub-types other than defined above are reserved for future extensions.

5. 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.

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.

6. Base MPLS Imposition MSD

The Base MPLS Imposition MSD (BMI-MSD) signals the total number of MPLS labels a node is capable of imposing, including any service/ transport labels.

Absence of BMI-MSD advertisements indicates solely that the advertising node does not support advertisement of this capability.

7. IANA Considerations

This document requests IANA to allocate TLV type (TBD1) from the OSPF Router Information (RI) TLVs Registry as defined by [RFC4970]. 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 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.

This document requests creation of an IANA managed registry under a new category of "Interior Gateway Protocol (IGP) Parameters" IANA registries to identify MSD types as proposed in Section 3, Section 4. The registration procedure is "Expert Review" as defined in [RFC8126]. The suggested registry name is "MSD types". 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	This document
251-254	Experimental	This document
255	Reserved	This document

Figure 3: MSD Types Codepoints Registry

8. Security Considerations

Security considerations, as specified by [RFC7770] are applicable to this document

9. Contributors

The following people contributed to this document:

Les Ginsberg

Email: ginsberg@cisco.com

10. Acknowledgements

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

11. References

11.1. 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>.

[RFC4970] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, DOI 10.17487/RFC4970, July 2007, <https://www.rfc-editor.org/info/rfc4970>.

[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>.

11.2. Informative References

[I-D.ietf-idr-bgp-ls-segment-routing-msd] Tantsura, J., Chunduri, U., Mirsky, G., and S. Sivabalan, "Signaling Maximum SID Depth using Border Gateway Protocol Link-State", draft-ietf-idr-bgp-ls-segment-routing-msd-01 (work in progress), October 2017.

[I-D.ietf-ospf-mpls-elc]
Xu, X., Kini, S., Sivabalan, S., Filsfils, C., and S. Litkowski, "Signaling Entropy Label Capability and Readable Label-stack Depth Using OSPF", draft-ietf-ospf-mpls-elc-05 (work in progress), January 2018.

[I-D.ietf-ospf-ospfv3-lsa-extend]
Lindem, A., Roy, A., Goethals, D., Vallem, V., and F. Baker, "OSPFv3 LSA Extendibility", draft-ietf-ospf-ospfv3-lsa-extend-23 (work in progress), January 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-11 (work in progress), November 2017.

[RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, April 2010, <https://www.rfc-editor.org/info/rfc5838>.

[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>.

[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>.

[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>.

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