

Internet Engineering Task Force (IETF)
Request for Comments: 9356
Updates: 9085
Category: Standards Track
ISSN: 2070-1721

K. Talaulikar, Ed.
P. Psenak
Cisco Systems
January 2023

Advertising Layer 2 Bundle Member Link Attributes in OSPF

Abstract

There are deployments where the Layer 3 (L3) interface on which OSPF operates is a Layer 2 (L2) interface bundle. Existing OSPF advertisements only support advertising link attributes of the L3 interface. If entities external to OSPF wish to control traffic flows on the individual physical links that comprise the L2 interface bundle, link attribute information for the bundle members is required.

This document defines the protocol extensions for OSPF to advertise the link attributes of L2 bundle members. The document also specifies the advertisement of these OSPF extensions via the Border Gateway Protocol - Link State (BGP-LS) and thereby updates RFC 9085.

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/rfc9356>.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
 - 1.1. Requirements Language
2. L2 Bundle Member Attributes
3. BGP-LS Advertisement
4. IANA Considerations
5. Operational Considerations
6. Security Considerations

7. References

7.1. Normative References

7.2. Informative References

Acknowledgements

Authors' Addresses

1. Introduction

There are deployments where the L3 interface on which an OSPF adjacency is established is a L2 interface bundle, for instance, a Link Aggregation Group (LAG) [IEEE802.1AX]. This reduces the number of adjacencies that need to be maintained by the OSPF protocol in cases where there are parallel links between the neighbors. Entities external to OSPF such as Path Computation Elements (PCEs) [RFC4655] may wish to control traffic flows on individual L2 member links of the underlying bundle interface (e.g., LAG). To do so, link attribute information for individual bundle members is required. The protocol extensions defined in this document provide the means to advertise this information.

This document defines sub-TLVs to advertise link attribute information for each of the L2 bundle members that comprise the L3 interface on which OSPF operates. Similar capabilities were introduced for IS-IS in [RFC8668].

[RFC8665] and [RFC8666] introduced the Adjacency Segment Identifier (Adj-SID) link attribute for OSPFv2 and OSPFv3, respectively, which can be used as an instruction to forward traffic over a specific link [RFC8402]. This document enables the advertisement of the Adj-SIDs using the same Adj-SID sub-TLV at the granularity level of each L2 bundle member link so that traffic may be steered over that specific member link.

Note that the advertisements at the L2 bundle member link level defined in this document are intended to be provided to entities external to OSPF and do not alter or change the OSPF route computation. The following items are intentionally not defined in and are outside the scope of this document:

- * What link attributes will be advertised. This is determined by the needs of the external entities.
- * A minimum or default set of link attributes.
- * How these attributes are configured.
- * How the advertisements are used.
- * What impact the use of these advertisements may have on traffic flow in the network.
- * How the advertisements are passed to external entities.

BGP Link State (BGP-LS) [RFC7752] was extended for the advertisement of L2 bundle members and their attributes in [RFC9085], which covered only IS-IS. This document updates [RFC9085] by specifying the advertisement from OSPF (refer to Section 3).

1.1. 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. L2 Bundle Member Attributes

A new L2 Bundle Member Attributes sub-TLV is introduced to advertise L2 bundle member attributes in both OSPFv2 and OSPFv3. In the case of OSPFv2, this sub-TLV is an optional sub-TLV of the OSPFv2 Extended Link TLV that is used to describe link attributes via the OSPFv2 Extended Link Opaque LSA (Link State Advertisement) [RFC7684]. In the case of OSPFv3, this sub-TLV is an optional sub-TLV of the Router-Link TLV of the OSPFv3 E-Router-LSA [RFC8362].

When the OSPF adjacency is associated with an L2 bundle interface, this sub-TLV is used to advertise the underlying L2 bundle member links along with their respective link attributes. The inclusion of this information implies that the identified link is a member of the L2 bundle associated with an OSPF L3 link and that the member link is operationally up. Therefore, advertisements of member links MUST NOT be done when the member link becomes operationally down or is no longer a member of the identified L2 bundle.

The advertisement of the L2 Bundle Member Attributes sub-TLV may be asymmetric for an OSPF link, depending on the underlying L2 connectivity, i.e., advertised by the router on only one end.

The L2 Bundle Member Attributes sub-TLV has the following format:

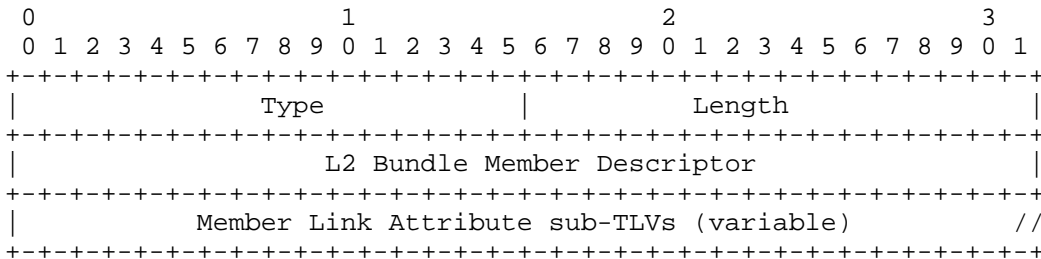


Figure 1: L2 Bundle Member Attributes Sub-TLV Format

Where:

Type: 24 for OSPFv2 and 29 for OSPFv3

Length: The total length (in octets) of the value portion of the TLV including nested sub-TLVs.

L2 Bundle Member Descriptor: A 4-octet link-local identifier for the member link. This identifier is described as "link local identifier" in [RFC4202] and used as "Local Interface ID" in [RFC8510].

Link attributes for L2 bundle member links are advertised as sub-TLVs of the L2 Bundle Member Attributes sub-TLV.

In the case of OSPFv2, the L2 Bundle Member Attributes sub-TLV shares the sub-TLV space of the Extended Link TLV, and the sub-TLVs of the Extended Link TLV MAY be used to describe the attributes of the member link. Table 1 lists sub-TLVs and their applicability for L2 bundle member links. The sub-TLVs that are not applicable MUST NOT be used as sub-TLVs for the L2 Bundle Member Attributes sub-TLV. Specifications that introduce new sub-TLVs of the Extended Link TLV MUST indicate their applicability to the L2 Bundle Member Attributes sub-TLV. Typically, attributes that have L3 semantics would not be applicable, but L2 attributes would apply. An implementation MUST ignore any sub-TLVs received that are not applicable in the context of the L2 Bundle Member Attributes sub-TLV.

=====+

Value	Description	Applicability
1	SID/Label	N
2	Adj-SID	Y
3	LAN Adj-SID/Label	Y
4	Network-to-Router Metric	N
5	RTM Capability	N
6	OSPFv2 Link MSD	N
7	Graceful-Link-Shutdown	N
8	Remote IPv4 Address	N
9	Local/Remote Interface ID	N
10	Application-Specific Link Attributes	Y
11	Shared Risk Link Group	Y
12	Unidirectional Link Delay	Y
13	Min/Max Unidirectional Link Delay	Y
14	Unidirectional Delay Variation	Y
15	Unidirectional Link Loss	Y
16	Unidirectional Residual Bandwidth	Y
17	Unidirectional Available Bandwidth	Y
18	Unidirectional Utilized Bandwidth	Y
19	Administrative Group	Y
20	Extended Administrative Group	Y
22	TE Metric	Y
23	Maximum Link Bandwidth	Y
24	L2 Bundle Member Attributes	N

Table 1: Applicability of OSPFv2 Link Attribute Sub-TLVs for L2 Bundle Members

Applicability:

Y: This sub-TLV MAY appear in the L2 Bundle Member Attributes sub-TLV.

N: This sub-TLV MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.

In the case of OSPFv3, the L2 Bundle Member Attributes sub-TLV shares the sub-TLV space of the Router-Link TLV, and the sub-TLVs of the Router-Link TLV MAY be used to describe the attributes of the member link. Table 2 lists sub-TLVs that are applicable to the Router-Link TLV and their applicability for L2 bundle member links. The sub-TLVs that are not applicable MUST NOT be used as sub-TLVs for the L2

Bundle Member Attributes sub-TLV. Specifications that introduce new sub-TLVs of the Router-Link TLV MUST indicate their applicability to the L2 Bundle Member Attributes sub-TLV. An implementation MUST ignore any sub-TLVs received that are not applicable in the context of the L2 Bundle Member Attributes sub-TLV.

Value	Description	Applicability
1	IPv6-Forwarding-Address	X
2	IPv4-Forwarding-Address	X
3	Route-Tag	X
4	Prefix SID	X
5	Adj-SID	Y
6	LAN Adj-SID	Y
7	SID/Label	N
8	Graceful-Link-Shutdown	N
9	OSPFv3 Link MSD	N
11	Application-Specific Link Attributes	Y
12	Shared Risk Link Group	Y
13	Unidirectional Link Delay	Y
14	Min/Max Unidirectional Link Delay	Y
15	Unidirectional Delay Variation	Y
16	Unidirectional Link Loss	Y
17	Unidirectional Residual Bandwidth	Y
18	Unidirectional Available Bandwidth	Y
19	Unidirectional Utilized Bandwidth	Y
20	Administrative Group	Y
21	Extended Administrative Group	Y
22	TE Metric	Y
23	Maximum Link Bandwidth	Y
24	Local Interface IPv6 Address	N
25	Remote Interface IPv6 Address	N
26	Flexible Algorithm Prefix Metric (FAPM)	X
27	Prefix Source OSPF Router-ID	X
28	Prefix Source Router Address	X
29	L2 Bundle Member Attributes	N
33	OSPF Flexible Algorithm ASBR Metric	X

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Table 2: Applicability of OSPFv3 Link Attribute Sub-TLVs for L2 Bundle Members

Applicability:

- Y: This sub-TLV MAY appear in the L2 Bundle Member Attributes sub-TLV.
- N: This sub-TLV MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.
- X: This is not a sub-TLV of the Router-Link TLV; it MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.

3. BGP-LS Advertisement

The BGP-LS extensions for the advertisement of L2 bundle members and their attributes were specified in [RFC9085]. Using the OSPF L2 Bundle Member Attributes sub-TLV defined in this document, the L2 bundle member information can now be advertised from OSPF into BGP-LS on the same lines as discussed for IS-IS in Section 2.2.3 of [RFC9085].

4. IANA Considerations

IANA has allocated the following code point in the "OSPFv2 Extended Link TLV Sub-TLVs" subregistry under the "Open Shortest Path First v2 (OSPFv2) Parameters" registry:

Value: 24

Designation: L2 Bundle Member Attributes

IANA has allocated the following code point in the "OSPFv3 Extended-LSA Sub-TLVs" subregistry under the "Open Shortest Path First v3 (OSPFv3) Parameters" registry:

Value: 29

Description: L2 Bundle Member Attributes

IANA has also introduced a column titled "L2BM" in the "OSPFv2 Extended Link TLV Sub-TLVs" registry. The "L2BM" column indicates applicability to the L2 Bundle Attributes Member sub-TLV. The initial allocations (Y/N) for this column are indicated in Table 1. The following explanatory note has been added to the registry:

The "L2BM" column indicates applicability to the L2 Bundle Attributes Member sub-TLV. The options for the "L2BM" column are:

Y - This sub-TLV MAY appear in the L2 Bundle Member Attributes sub-TLV.

N - This sub-TLV MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.

Similarly, IANA has introduced a column titled "L2BM" in the "OSPFv3 Extended-LSA Sub-TLVs" registry. The "L2BM" column indicates applicability to the L2 Bundle Attributes Member sub-TLV. The initial allocations (Y/N/X) for this column are indicated in Table 2. The following explanatory note has been added to the registry:

The "L2BM" column indicates applicability to the L2 Bundle Attributes Member sub-TLV. The options for the "L2BM" column are:

Y - This sub-TLV MAY appear in the L2 Bundle Member Attributes sub-TLV.

N - This sub-TLV MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.

X - This is not a sub-TLV of the Router-Link TLV; it MUST NOT appear in the L2 Bundle Member Attributes sub-TLV.

Future allocations in these two registries are required to indicate the applicability of the introduced sub-TLV to the L2 Bundle Member Attributes sub-TLV. IANA has added this document as a reference for both registries.

5. Operational Considerations

Implementations MUST NOT enable the advertisement of L2 bundle member links and their attributes in OSPF LSAs by default and MUST provide a configuration option to enable their advertisement on specific links.

[RFC9129] specifies the base YANG data model for OSPF. The required configuration and operational elements for this feature are expected to be introduced as augmentation to this base YANG data model for OSPF.

6. Security Considerations

The OSPF protocol has supported the advertisement of link attribute information, including link identifiers, for many years. The advertisements defined in this document are identical to the existing advertisements defined in [RFC3630], [RFC4203], [RFC5329], [RFC7471], [RFC8665], and [RFC8666], but they are associated with L2 links that are part of a bundle interface on which the OSPF protocol operates. Therefore, the security considerations of these documents are applicable, and there are no new security issues introduced by the extensions in this document.

As always, if the protocol is used in an environment where unauthorized access to the physical links on which OSPF packets are sent occurs, then attacks are possible. The use of authentication as defined in [RFC5709], [RFC7474], [RFC4552], and [RFC7166] is recommended for preventing such attacks.

7. References

7.1. Normative References

[IEEE802.1AX]

IEEE, "IEEE Standard for Local and Metropolitan Area Networks--Link Aggregation", IEEE Std 802.1AX, DOI 10.1109/IEEESTD.2020.9105034, May 2020, <<https://doi.org/10.1109/IEEESTD.2020.9105034>>.

[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>>.

[RFC4202] Kompella, K., Ed. and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, DOI 10.17487/RFC4202, October 2005, <<https://www.rfc-editor.org/info/rfc4202>>.

[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>>.

- [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>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", RFC 8666, DOI 10.17487/RFC8666, December 2019, <<https://www.rfc-editor.org/info/rfc8666>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.

7.2. Informative References

- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<https://www.rfc-editor.org/info/rfc4203>>.
- [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>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <<https://www.rfc-editor.org/info/rfc5329>>.
- [RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", RFC 5709, DOI 10.17487/RFC5709, October 2009, <<https://www.rfc-editor.org/info/rfc5709>>.
- [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>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric

Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015,
<<https://www.rfc-editor.org/info/rfc7471>>.

[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>>.

[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>>.

[RFC8510] Psenak, P., Ed., Talaulikar, K., Henderickx, W., and P.
Pillay-Esnault, "OSPF Link-Local Signaling (LLS)
Extensions for Local Interface ID Advertisement",
RFC 8510, DOI 10.17487/RFC8510, January 2019,
<<https://www.rfc-editor.org/info/rfc8510>>.

[RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri,
M., and E. Aries, "Advertising Layer 2 Bundle Member Link
Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668,
December 2019, <<https://www.rfc-editor.org/info/rfc8668>>.

[RFC9129] Yeung, D., Qu, Y., Zhang, Z., Chen, I., and A. Lindem,
"YANG Data Model for the OSPF Protocol", RFC 9129,
DOI 10.17487/RFC9129, October 2022,
<<https://www.rfc-editor.org/info/rfc9129>>.

Acknowledgements

This document leverages similar work done for IS-IS, and the authors
of this document would like to acknowledge the contributions of the
authors of [RFC8668].

The authors would like to thank Anoop Ghanwani, Paul Kyzivat, Dan
Romascanu, and Russ Mundy for their review and feedback on this
document. The authors would also like to thank Acee Lindem for his
detailed shepherd review of this document. The authors would also
like to thank John Scudder for his AD review and the discussion
related to the applicability of TLVs/sub-TLVs to the L2 Bundle Member
Attributes sub-TLV.

Authors' Addresses

Ketan Talaulikar (editor)
Cisco Systems
India
Email: ketant.ietf@gmail.com

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynske nivy 43
821 09 Bratislava
Slovakia
Email: ppsenak@cisco.com