

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 19, 2016

L. Ginsberg
A. Bashandy
C. Filsfils
S. Previdi
Cisco Systems
M. Nanduri
Microsoft
E. Aries
Private Contributor
February 16, 2016

**Advertising L2 Bundle Member Link Attributes in IS-IS
draft-ginsberg-isis-l2bundles-02.txt**

Abstract

There are deployments where the Layer 3 interface on which IS-IS operates is a Layer 2 interface bundle. Existing IS-IS advertisements only support advertising link attributes of the Layer 3 interface. If entities external to IS-IS wish to control traffic flows on the individual physical links which comprise the Layer 2 interface bundle link attribute information about the bundle members is required.

This document introduces the ability for IS-IS to advertise the link attributes of layer 2 (L2) bundle members.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 <http://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 August 19, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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](#)
- [2.](#) L2 Bundle Member Attributes TLV [3](#)
 - [2.1.](#) Parallel L3 Adjacencies [5](#)
 - [2.2.](#) Shared Attribute sub-TLVs [5](#)
- [3.](#) Advertising L2 Bundle Member Adj-SIDs [5](#)
 - [3.1.](#) L2 Bundle Member Adjacency Segment Identifier sub-TLV . . [6](#)
 - [3.2.](#) L2 Bundle Member LAN Adjacency Segment Identifier sub-TLV [7](#)
- [4.](#) IANA Considerations [9](#)
- [5.](#) Security Considerations [11](#)
- [6.](#) Acknowledgements [11](#)
- [7.](#) References [11](#)
 - [7.1.](#) Normative References [11](#)
 - [7.2.](#) Informational References [12](#)
- Authors' Addresses [12](#)

[1.](#) Introduction

There are deployments where the Layer 3 interface on which an IS-IS adjacency is established is a Layer 2 interface bundle, for instance a Link Aggregation Group (LAG) [[IEEE802.1AX](#)]. This reduces the number of adjacencies which need to be maintained by the routing protocol in cases where there are parallel links between the neighbors. Entities external to IS-IS such as Path Computation Elements (PCE) [[RFC4655](#)] may wish to control traffic flows on individual members of the underlying Layer 2 bundle. In order to do so link attribute information about individual bundle members is required - but currently IS-IS only supports advertising link attributes for the Layer 3 interfaces on which it operates.

This document introduces a new TLV to advertise link attribute information for each of the L2 bundle members which comprise the Layer 3 interface on which IS-IS operates.

[SR] introduces a new link attribute - adjacency segment identifier (Adj-SID) - which can be used as an instruction to forwarding to send traffic over a specific link. This document introduces additional sub-TLVs to advertise Adj-SIDs for L2 Bundle members.

Note that the new advertisements defined in this document are intended to be provided to external entities.

2. L2 Bundle Member Attributes TLV

A new TLV is introduced to advertise L2 Bundle member attributes. Although much of the information is identical to and uses the same sub-TLVs included in Extended IS-Neighbor advertisements (TLVs 22 and 222), a new TLV is used so that changes to the advertisement of the L2 Bundle member link attributes does not trigger unnecessary action by the [[ISO10589](#)] Decision process.

This new TLV utilizes the sub-TLV space defined for TLVs 22, 23, 141, 222, and 223.

The following new TLV is introduced:

L2 Bundle Member Attributes


Type: 25 (suggested - to be assigned by IANA)

Length: Number of octets to follow

Parent L3 Neighbor Descriptor

L3 Neighbor System ID + pseudonode ID (7 octets)

Flags: 1 octet field of following flags:



where:

P-flag: When set to 1 one of the sub-TLVs described in [Section 2.1](#) immediately follows the flags field. If the P-flag is set to 0, then none of the sub-TLVs described in [Section 2.1](#) are present.

Other bits: MUST be zero when originated and ignored when received.

One or more of the following:

L2 Bundle Attribute Descriptors

Length of L2 Bundle Attribute Descriptor (1 octet)

NOTE: This includes all fields described below.

Number of L2 Bundle Member Descriptors (1 octet)

L2 Bundle Member Link Local Identifiers

(4 * Number of L2 Bundle Member Descriptors octets)

NOTE: An L2 Bundle Member Descriptor is a Link Local Identifier as defined in [\[RFC5307\]](#).

sub-TLV(s)

A sub-TLV may define an attribute common to all of the bundle members listed or a sub-TLV may define an attribute unique to each bundle member. Use of these two classes of sub-TLVs is described in the following sections.

NOTE: Only one Parent L3 Neighbor Descriptor is present in a given TLV. Multiple L2 Bundle Attribute Descriptors may be present in a single TLV.

2.1. Parallel L3 Adjacencies

When there exist multiple L3 adjacencies to the same neighbor additional information is required to uniquely identify the L3 Neighbor. One and only one of the following three sub-TLVs is used to uniquely identify the L3 adjacency:

- o IPv4 Interface Address (sub-TLV 6 defined in [[RFC5305](#)])
- o IPv6 Interface Address (sub-TLV 12 defined in [[RFC6119](#)])
- o Link Local/Remote Identifiers (sub-TLV 4 defined in [[RFC5307](#)])

When the P-bit is set in the flags field in the Parent L3 Neighbor Descriptor one and only one of the above sub-TLVs MUST be present. The chosen sub-TLV MUST immediately follow the flags field described in [Section 2](#).

These sub-TLVs MAY be omitted if no parallel adjacencies to the neighbor exist.

2.2. Shared Attribute sub-TLVs

These sub-TLVs advertise a single copy of an attribute (e.g. link bandwidth). The attribute applies to all of the L2 Bundle Members in the set advertised under the preceding L2 Bundle Member Attribute Descriptor. No more than one copy of a given sub-TLV in this category may appear in the set of sub-TLVs under the preceding L2 Bundle Member Attribute Descriptor. If multiple copies of a given sub-TLV are present both MUST be ignored.

The set of L2 Bundle Member Descriptors which may be advertised under a single L2 Bundle Member Attribute Descriptor is therefore limited to bundle members which share the set of attributes advertised in the shared attribute sub-TLVs.

All existing sub-TLVs defined in the IANA Sub-TLVs for TLVs 22, 23, 141, 222, and 223 registry are in the category of shared attribute sub-TLVs unless otherwise specified in this document.

3. Advertising L2 Bundle Member Adj-SIDs

[SR] defines sub-TLVs to advertise Adj-SIDs for L3 adjacencies. However these sub-TLVs only support an advertisement of a single Adj-SID. As it is expected that each L2 Bundle member will have unique Adj-SIDs in many deployments it is desirable to define a new sub-TLV which allows more efficient encoding of a set of Adj-SIDs in a single sub-TLV. Two new sub-TLVs are therefore introduced to support

advertising Adj-SIDs for L2 Bundle members. The format of the new sub-TLVs is similar to that used for L3 adjacencies, but is optimized to allow advertisement of a set of Adj-SIDs (one per L2 Bundle Member) in a single sub-TLV.

The two new sub-TLVs defined in the following sections do not fall into the category of shared attribute sub-TLVs.

3.1. L2 Bundle Member Adjacency Segment Identifier sub-TLV

This sub-TLV is used to advertise Adj-SIDs for L2 Bundle Members associated with a parent L3 adjacency which is Point-to-Point. The following format is defined for this sub-TLV:

Type: 41 (suggested value to be assigned by IANA) (1 octet)
Length: variable (1 octet)

Flags: 1 octet field of following flags:

```
  0 1 2 3 4 5 6 7
  +--+--+--+--+--+
  |F|*|V|L|S|  |
  +--+--+--+--+--+
```

where:

* - Is a flag used in the L3 Adj-SID sub-TLV but which is NOT used in this sub-TLV. These bits SHOULD be sent as 0 and MUST be ignored on receipt

F-Flag: Address-Family flag. If unset, then the Adj-SID refers to an L2 Bundle Member with outgoing IPV4 encapsulation. If set then the Adj-SID refers to an L2 Bundle Member with outgoing IPV6 encapsulation.

V-Flag: Value flag. If set, then the Adj-SID carries a value. By default the flag is SET.

L-Flag: Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.

S-Flag. Set Flag. When set, the S-Flag indicates that the Adj-SID refers to a set of L2 Bundle Members (and therefore MAY be assigned to other L2 Bundle Members as well).

Other bits: MUST be zero when originated and ignored when

received.

Weight: 1 octet. The value represents the weight of the Adj-SID for the purpose of load balancing. The use of the weight is defined in [[SR-ARCH](#)].

NOTE: Flags and weight are shared by all L2 Bundle Members listed in the L2 Bundle Attribute Descriptor.

L2 Bundle Member Adj-SID Descriptors. There MUST be one descriptor for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor. Each descriptor consists of one of the following fields:

SID/Index/Label: according to the V and L flags, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the SID/Label space advertised by this router. See [[SR](#)]. In this case V and L flags MUST be unset.
- * A 16 octet IPv6 address. In this case the V flag MUST be set. The L flag MUST be unset if the IPv6 address is globally unique.

3.2. L2 Bundle Member LAN Adjacency Segment Identifier sub-TLV

This sub-TLV is used to advertise Adj-SIDs for L2 Bundle Members associated with a parent L3 adjacency which is a LAN adjacency. In LAN subnetworks, the Designated Intermediate System (DIS) is elected and originates the Pseudonode-LSP (PN-LSP) including all neighbors of the DIS. When Segment Routing is used, each router in the LAN MAY advertise the Adj-SID of each of its neighbors on the LAN. Similarly, for each L2 Bundle Member a router MAY advertise an Adj-SID to each neighbor on the LAN.

The following format is defined for this sub-TLV:

Type: 42 (suggested value to be assigned by IANA) (1 octet)
Length: variable (1 octet)
Neighbor System ID: 6 octets

Flags: 1 octet field of following flags:

```

 0 1 2 3 4 5 6 7
+-+--+--+--+--+
|F|*|V|L|S|  |
+-+--+--+--+--+
    
```

where:

* - Is a flag used in the L3 Adj-SID sub-TLV but which is NOT used in this sub-TLV. These bits SHOULD be sent as 0 and MUST be ignored on receipt

F-Flag: Address-Family flag. If unset, then the Adj-SID refers to an L2 Bundle Member with outgoing IPv4 encapsulation. If set then the Adj-SID refers to an L2 Bundle Member with outgoing IPv6 encapsulation.

V-Flag: Value flag. If set, then the Adj-SID carries a value. By default the flag is SET.

L-Flag: Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.

S-Flag. Set Flag. When set, the S-Flag indicates that the Adj-SID refers to a set of L2 Bundle Members (and therefore MAY be assigned to other L2 Bundle Members as well).

Other bits: MUST be zero when originated and ignored when received.

Weight: 1 octet. The value represents the weight of the Adj-SID for the purpose of load balancing. The use of the weight is defined in [[SR-ARCH](#)].

NOTE: Flags and weight are shared by all L2 Bundle Members listed in the L2 Bundle Attribute Descriptor.

L2 Bundle Member LAN Adj-SID Descriptors. There MUST be one descriptor for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor. Each descriptor consists of one of the following fields:

SID/Index/Label: according to the V and L flags, it contains either:

* A 3 octet local label where the 20 rightmost bits are used

for encoding the label value. In this case the V and L flags MUST be set.

- * A 4 octet index defining the offset in the SID/Label space advertised by this router. See [\[SR\]](#). In this case V and L flags MUST be unset.
- * A 16 octet IPv6 address. In this case the V flag MUST be set. The L flag MUST be unset if the IPv6 address is globally unique.

4. IANA Considerations

This document adds the following new TLV to the IS-IS TLV Codepoints registry.

Value: 25 (suggested - to be assigned by IANA)

Name: L2 Bundle Member Attributes

The name of the Sub-TLVs for TLVs 22, 23, 141, 222, and 223 registry needs to be changed to Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 registry. An additional column needs to be added to the registry to indicate which sub-TLVs may appear in the new L2 Bundle Member Attributes TLV. The following table indicates the appropriate settings for all currently defined sub-TLVs as regards their use in the new L2 Bundle Member Attributes TLV.

3 Administrative group (color) y
 4 Link Local/Remote Identifiers y
 6 IPv4 interface address y
 8 IPv4 neighbor address y
 9 Maximum link bandwidth y
 10 Maximum reservable link bandwidth y
 11 Unreserved bandwidth y
 12 IPv6 Interface Address y
 13 IPv6 Neighbor Address y
 14 Extended Administrative Group y
 18 TE Default metric y
 19 Link-attributes y
 20 Link Protection Type y
 21 Interface Switching Capability Descriptor y
 22 Bandwidth Constraints y
 23 Unconstrained TE LSP Count y
 24 Remote AS number n
 25 IPv4 remote ASBR Identifier n
 26 IPv6 remote ASBR Identifier n
 27 Interface Adjustment Capability Descriptor (IACD) y
 28 MTU n
 29 SPB-Metric y
 30 SPB-A-OALG y

This document adds the following new sub-TLVs to the sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 registry.

Value: 41 (suggested - to be assigned by IANA)

Name: L2 Bundle Member Adj-SID

This sub-TLV is allowed in the following TLVs:

| | | | | | |
|----|----|----|-----|-----|-----|
| 22 | 23 | 25 | 141 | 222 | 223 |
| n | n | y | n | n | n |

Value: 42 (suggested to be assigned by IANA)

Name: L2 Bundle Member LAN Adj-SID

This sub-TLV is allowed in the following TLVs:

| | | | | | |
|----|----|----|-----|-----|-----|
| 22 | 23 | 25 | 141 | 222 | 223 |
| n | n | y | n | n | n |

5. Security Considerations

None.

6. Acknowledgements

The authors would like to thank Jon Mitchell for his careful review.

7. References

7.1. Normative References

[IEEE802.1AX]

Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation.", ISO/IEC 10589:2002, Second Edition, Nov 2008.

[ISO10589]

International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, Nov 2002.

[RFC2119]

Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC5305]

Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.

[RFC5307]

Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 5307](#), DOI 10.17487/RFC5307, October 2008, <<http://www.rfc-editor.org/info/rfc5307>>.

[RFC6119]

Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", [RFC 6119](#), DOI 10.17487/RFC6119, February 2011, <<http://www.rfc-editor.org/info/rfc6119>>.

7.2. Informational References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [SR] "IS-IS Extensions for Segment Routing, [draft-ietf-isis-segment-routing-extensions-06](#)(work in progress)", December 2015.
- [SR-ARCH] "Segment Routing Architecture, [draft-ietf-spring-segment-routing-07](#)(work in progress)", December 2015.

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: ginsberg@cisco.com

Ahmed Bashandy
Cisco Systems
170 West Tasman Drive
San Jose, Ca 95134
US

Clarence Filsfils
Cisco Systems

Email: cf@cisco.com

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Mohan Nanduri
Microsoft

Email: mmanduri@microsoft.com

Ebben Aries
Private Contributor

Email: exa@dscp.org