

Advertising L2 Bundle Member Link Attributes in IS- IS

draft-ginsberg-isis-l2bundles-00.txt

Les Ginsberg (ginsberg@cisco.com)

Ahmed Bashandy (bashandy@cisco.com)

Clarence Filsfils(cfilsfils@cisco.com)

Stefano Previdi (sprevidi@cisco.com)

Mohan Nanduri (mnanduri@microsoft.com)

Ebben Aries (exa@fb.com)

Motivation

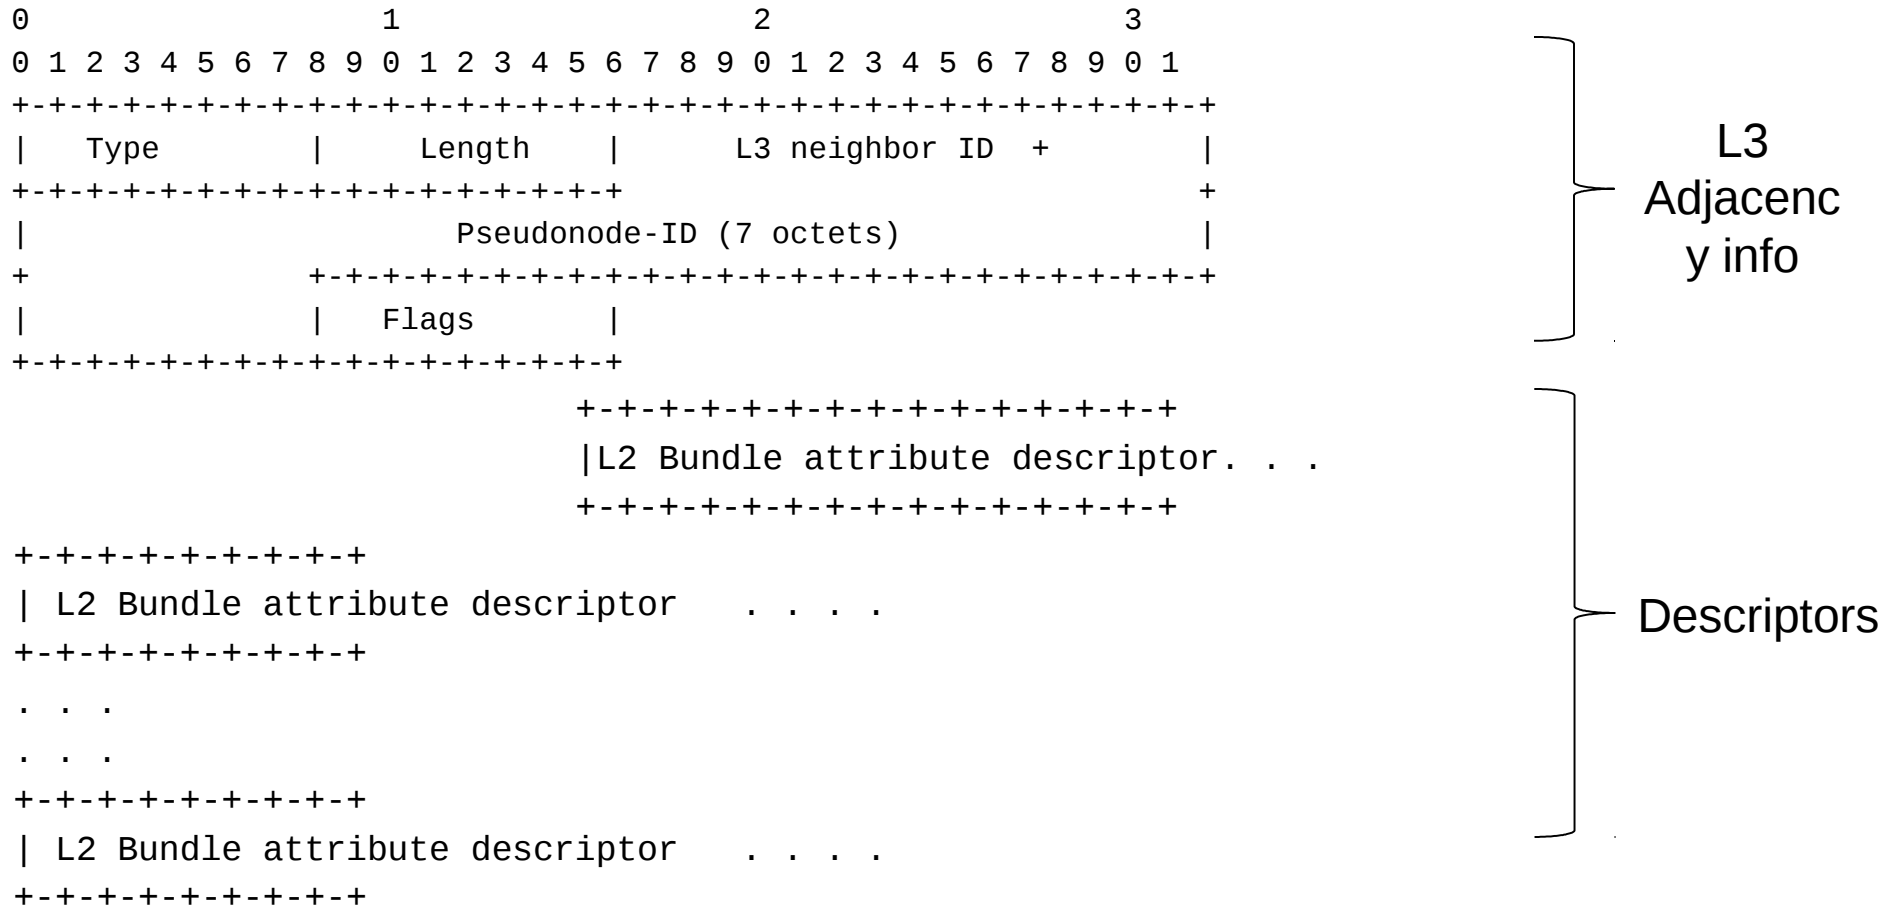
- Wide spread use of L2 link aggregation
 - Bandwidth gain while maintaining cost effectiveness
 - Increase reliability
 - Scalability: Shield L3 protocols from possibly large number of links
 - Provisioning simplicity: 1 IP prefix for all member links
- What we want to do ?
 - Individual member health monitor
 - Traffic engineering over individual members
- Scalability concern:
 - Do *not* want to expose the possibly very large number of links to all protocols

L2 bundle member Attributes

TLV

- Top level TLV
- Why top level TLV?
 - Isolates L3 topology changes from L2 attributes
 - Not intermixed (potentially) with L3 link attributes, which may have simpler syntax
- Advertises individual links and their attributes
- Shares same sub-TLV space with Extended IS reachability TLVs (22, 23, 141, 222, 223)
 - We will modify the IANA registries for these sub-TLVs
- The TLV consists of 2 main parts
 - Parent L3 adjacency information
 - One or more L2 bundle Attribute Descriptor

L2 bundle member Attributes TLV



Type: 25 (Suggested value to be assigned by IANA)

Length: Length of the TLV

Parent L3 Adjacency Information

```

      0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type          |  Length      |  L3 neighbor ID  +  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Pseudonode-ID (7 octets) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |  Flags          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

- **L3 Neighbor System ID + pseudonode ID** (7 octets)
 - Identifies the neighbor
- **Flags:** 1 octet field of the following flags

```

      0 1 2 3 4 5 6 7
+--+--+--+--+--+--+--+
|P|          |
+--+--+--+--+--+--+--+

```

P flag

If set, then one of the following **MUST** immediately follow the **flags** field

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

The “P” flag **MUST** be set for *parallel* adjacencies

L2 Bundle Attribute Descriptor

- Describes a (sub-)set of the members of the L2 bundle
 - Includes sub-TLVs that describe attributes applicable to that set of members

- Structure

L2 Bundle Attribute descriptor Length(1 octet)

Includes all fields to follow including sub-TLVs

Number of L2 bundle members: (1 octet)

L2 Bundle members IDs (4 * Number of L2 bundle members)

Zero or more sub-TLVs

- The sub-TLVs to follow specify attributes applicable to all members listed in this descriptor

L2 Bundle Attribute Descriptor

- Two types of attributes
 - **Shared attributes**: All sub-TLVs used in extended reachability TLV define attributes applicable to all members listed in the descriptor
 - **Individual member attributes**: A sub-TLV that specifies different attributes for different members
 - We define **two sub-TLVs** that specify attributes describing individual member properties
- Example of shared attributes:
 - Suppose maximum link bandwidth sub-TLV (sub-TLV 9 in RFC5305) exists in the descriptor
 - Then all members listed the descriptor share the same maximum link bandwidth
- Example of Individual member sub-TLVs
 - P2P and LAN L2 bundle member adjacency SIDs
 - These are the two only **new** sub-TLVs defined in this draft

L2 Bundle Member P2P adj-SID

- Specifies an adjacency SID per bundle member:
- Structure

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type          |      Length      |      Flags      |      weight      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  SID/Index/Label (variable, depends on size, V and L) . . .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
. . . .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  SID/Index/Label (variable, depends on size, V and L) . . .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
. . . .
. . . .

```


Type: Suggested 41 (suggested value to be assigned by IANA)

Length: Length of the sub-TLV

L2 Bundle Member P2P Adj-SID

(2)

- Flags: 1 octet field of following flags:



* - a flag used in the L3 Adj-SID sub-TLV (draft-ietf-isis-segment-routing-extensions) but NOT used in this sub-TLV. SHOULD be sent as 0 and MUST ignored on receipt

F-Flag: Address-Family flag.

Clear → Adj-SID refers outgoing IPv4 encapsulation.

Set → Adj-SID refers to outgoing IPv6 encapsulation.

L-Flag: Local Flag. *Set* → Adj-SID is local significance. By default the flag is SET.

S-Flag: Set Flag. *Set* → Adj-SID refers to a set of Members

L2 Bundle Member P2P Adj-SID

(2)

- **Weight (1 octet):** The weight of the adj-SIDs. See (draft-ietf-spring-segment-routing)
- **Adj-SIDs**
 - There **MUST** be one **SID/Index/Label** per bundle member advertised in the enclosing *L2 Bundle Member Attribute*
- **SID/Index/Label:**
 - V and L flag *set* and size = 2 + number of members * 3 :
 - 3 octet local label encoded in the 20 rightmost bits
 - V and L flag *Clear* and size = 2 + number of members * 4 :
 - 4 octet specifying the index from the SRGB (draft-ietf-isis-segment-routing-extensions)
 - V *set* and size = 2 + number of members * 16 :
 - 16 octet IPv6 address
 - If “L” flag is *clear*, then the IPv6 address is *globally* unique

L2 Bundle Member LAN Adj-SID

- Specifies an adj-SID per bundle member per ISIS neighbor
 - MUST use a *separate descriptor* for each ISIS neighbor
- Type: 42 (Suggested value to be assigned by IANA)
- Structure

```

      0                1                2                3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type          |  Length      |  Flags      |  weight     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                Neighbor System-ID (6 octets)                |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                     +--+--+--+--+--+--+--+--+--+--+
                                     |  SID/index/label    . . .
                                     +--+--+--+--+--+--+--+--+--+--+

+--+--+--+--+--+--+--+--+--+--+
|sid/index/label  . . . .
+--+--+--+--+--+--+--+--+--+--+

                                     +--+--+--+--+--+--+--+--+--+--+
                                     |  SID/index/label    . . .
                                     +--+--+--+--+--+--+--+--+--+--+

+--+--+--+--+--+--+--+--+--+--+
|sid/index/label  . . . .
+--+--+--+--+--+--+--+--+--+--+

. . . .
. . . .

```

Remaining fields are the same as P2P but towards the specific neighbor

Example


- Suppose a P2P L2 bundle has 4 members where 2 members are 10G and 2 members are 1G
- Suppose the advertising router allocates the local MPLS label values 24000, 24001, 24002, and 24003 for adjacency SIDs to the members.
- Suppose the router wants to advertise the maximum bandwidth sub-TLV (sub-TLV 9) for the members
- In this case we will have 2 separate descriptors
 - First descriptor for the 10G members
 - Second descriptor for the 1G members

Example: The over all TLV

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Type = 25   | Length = 60   |      L3 neighbor ID   +   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Pseudonode-ID (7 octets) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |0|                       |
+-----+-----+-----+-----+-----+-----+-----+-----+
                                     +-----+-----+-----+-----+
                                     |First Descriptor . . .
                                     +-----+-----+-----+-----+
. . . . .
. . . . .
+-----+-----+-----+-----+-----+-----+-----+-----+
| End of First descriptor |
+-----+-----+-----+-----+-----+-----+-----+-----+
                                     +-----+-----+-----+-----+
                                     |Second Descriptor . . .
                                     +-----+-----+-----+-----+
. . . . .
. . . . .
+-----+-----+-----+-----+-----+-----+-----+-----+
| End of Second descriptor|
+-----+-----+-----+-----+-----+-----+-----+-----+

```



Example(cont): First Descriptor

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| length = 25   | num-members = 2 | ID1 = |
+-----+-----+-----+-----+
|           0x100001           |
+-----+-----+-----+-----+
|                                     |
|                                     | ID2 = |
|                                     |
+-----+-----+-----+-----+
|           0X100002           |
+-----+-----+-----+-----+
|                                     |
|                                     | TYPE= 41 | Length = 8 |
+-----+-----+-----+-----+
|0|0|1|1|0|0|0|0| Weight = 10 | SID/Index/Label = |
+-----+-----+-----+-----+
| 24000 | SID/Index/Label = 24001 |
+-----+-----+-----+-----+
|                                     |
| Type = 9   | Length = 4   | Max Bandwidth = |
+-----+-----+-----+-----+
|           1.25 Mbytes/second           |
+-----+-----+-----+-----+

```

Members included in
the First descriptor

L2 bundle adj-SID sub-TLV

Max BW sub-TLV

Example(cont): Second Descriptor

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| length = 25   | num-members = 2 | ID3 = |
+-----+-----+-----+-----+-----+-----+-----+
|               | 0x100003       |
+-----+-----+-----+-----+-----+-----+
|               | ID4 =         |
+-----+-----+-----+-----+-----+-----+
|               | 0X100004     |
+-----+-----+-----+-----+-----+-----+
|               | TYPE= 41   | Length = 8 |
+-----+-----+-----+-----+-----+-----+
|0|0|1|1|0|0|0|0| Weight = 1 | SID/Index/Label = |
+-----+-----+-----+-----+-----+-----+
| 24002 | SID/Index/Label = 24003 |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Type = 9   | Length = 4   | Max Bandwidth = |
+-----+-----+-----+-----+-----+-----+
| 0.125 Mbytes/second |
+-----+-----+-----+-----+-----+-----+

```

} Members included in the Second descriptor

} L2 bundle adj-SID sub-TLV

} Max BW sub-TLV

Request to become a Working Group Document

!!

