

# Update on Advertising L2 Bundle Member Link Attributes in IS-IS

draft-ginsberg-isis-l2bundles-00.txt

Les Ginsberg ([ginsberg@cisco.com](mailto:ginsberg@cisco.com))

Ahmed Bashandy ([bashandy@cisco.com](mailto:bashandy@cisco.com))

Clarence Filsfils([cfilsfils@cisco.com](mailto:cfilsfils@cisco.com))

Stefano Previdi ([sprevidi@cisco.com](mailto:sprevidi@cisco.com))

Mohan Nanduri ([mnanduri@microsoft.com](mailto:mnanduri@microsoft.com))

Ebben Aries ([exa@fb.com](mailto:exa@fb.com))

# Status

- So far very minor comments on the draft itself
  - Related to references to the IEEE LAG standard
- Some discussion on why the draft is needed
- Few alternative proposals

## Discussion: Why the draft is needed

- Already mentioned in the draft and in the last IETF (IETF-93 in Prague)
- Need TE on individual member links
- Need to steer traffic into individual member links
- Do not want to overwhelm all routers and management tools with significant increase in number of links
- Do not want to force user to change the deployment strategy from L2 bundles to L3 links

# Discussion: Alternative Proposals

- Alternative proposals on the mailing list
  - Use of BGP-LS
  - Use unnumbered interface instead of bundles
  - IP address on individual bundle members as well as the bundle itself
- ) Common point in all 3 alternative proposals
  - Want to expose the L2 members to layer 3

## (1) Use of BGP-LS

- Proposal is not clear
- We do not understand how BGP-LS can be used as an alternative.  
E.g.
  - Deploy BGP on all nodes in the core?
  - Use BGP-LS to exchange link states instead of ISIS/OSPF?
- We think BGP-LS is *complimentary* not alternative


# Discussion: Alternative Proposals

## (2) Use unnumbered interface instead of bundles

- Instead of configuring a bundle, configure each member as an L3 link
- Make all members unnumbered (e.g. to Loopback0)

## (3) IP address on the individual members as well as the bundle itself

- Example sent on the list



```
interface Port-Channel1
 ip address 1.1.1.1/30
 ip router isis
!
interface Eth0
 ip address 2.2.2.1/30
 ip router isis
 channel-group 1
!
interface Eth1
 ip address 2.2.3.1/30
 ip router isis
 channel-group 1
```

## Comparing Proposals

	<b>ginsberg-isis-l2bundles</b>	<b>BGP-LS advertise L2 bundle</b>	<b>L2 bundles as unnumbered interfaces</b>	<b>Mix of L2 and L3 bundles as in ISIS</b>
Scalability	Minimal scale overhead (+1)	Minimal scale overhead (+1)	Significant scale overhead: Number of advertised links and ISIS adjacencies are multiplied by average number of bundle members(-1)	VERY significant Scale overhead - # of advertised links and ISIS adjacencies multiplied by number of bundle members - # of advertised and downloaded prefixes multiplied by the average number of bundle members (-2)
Deploy protocol not deployed before	No (+1)	Yes (-1)	No (+1)	No (+1)
Impact on basic routing functionality	Minimal (+1)	Minimal (+1)	Significant (-1): Sharp increase in the number of links in LSDB and SPF	VERY Significant (-2): - Sharp increase in the number of links in LSDB and SPF - sharp increase in number of prefixes to download to RIB and FIB
Works for both P2P and LAN	Simple (+1)	Simple (+1)	Difficult (-1)	Simple (+1)
Using 1 protocol for diverse functionalities	Yes (+1)	No (-1)	Yes (+1)	Yes (+1)
Exposing L2 info in L3 protocol	Yes (-1)	Yes (-1)	No (+1)*	No (+1)*

## Comparing Proposals

	<b>ginsberg-isis-l2bundles</b>	<b>BGP-LS advertise L2 bundle</b>	<b>L2 bundles as unnumbered interfaces</b>	<b>Mix of L2 and L3 bundles as in ISIS</b>
Protocol change	Yes (-1)	Yes (-1)	No (+1)	No (+1)
Risk	Small (Minimal impact on baseline functionality and managements tools while not deploying new protocol) (+1)	Medium (Have to deploy BGP-LS everywhere) (0)	Significant (MUST make sure that baseline functionality on all routers as well as management and monitoring tools are not impacted by the sharp increase in number of links) (-1)	VERY Significant (Have to make sure that baseline functionality on all routers as well as management and monitoring tools are not impacted by the sharp scale increase in both number of links and IP prefixes) (-2)
Provisioning Overhead	Small: Only need to enable the CLI under ISIS (+1)	Medium: (Have to configure BGP-LS everywhere) (0)	Significant: Have to deconfigure the bundle interfaces and make members unnumbered then add the member interfaces under ISIS process (-1)	VERY significant: Have to configure and maintain a very large number of prefixes between adjacent routers and make sure there is no prefix mismatch between peers (-2)
Code change	Medium: Only change ISIS (0)	Medium: Only change BGP (0)	None: can be achieved with current code (+1)	Very intrusive code change (-2) - MUST make changes to many applications on router to selectively ignore or handle significant increase in number of prefixes and L3 links - Have to make changes to the management and monitoring tools to selectively ignore and/or handle increase in both number links and number of prefixes

## Comparing Proposals

	<b>ginsberg-isis-l2bundles</b>	<b>BGP-LS advertise L2 bundle</b>	<b>L2 bundles as unnumbered interfaces</b>	<b>Mix of L2 and L3 bundles as in ISIS</b>
Works with RSVP without RSVP modification	No (-1)	No (-1)	Yes(+1)	Yes (+1)
Two way connectivity check	Simple because the draft allows advertising most of subTLVs in TLV22 in the bundle attribute TLV. Hence a router can include Sub-TLV 4 (+1)*	I do not think it is possible (-1)	Simple using existing ISIS functionality (+1)	Simple using existing ISIS functionality (+1)
Must deconfigure/reconfigure L2 bundles	No (+1)	No (+1)	yes (-1)	Yes (-1)
<b>Sum</b>	<b>6</b>	<b>-2</b>	<b>1</b>	<b>-3</b>



# Request to become a Working Group Document

!!



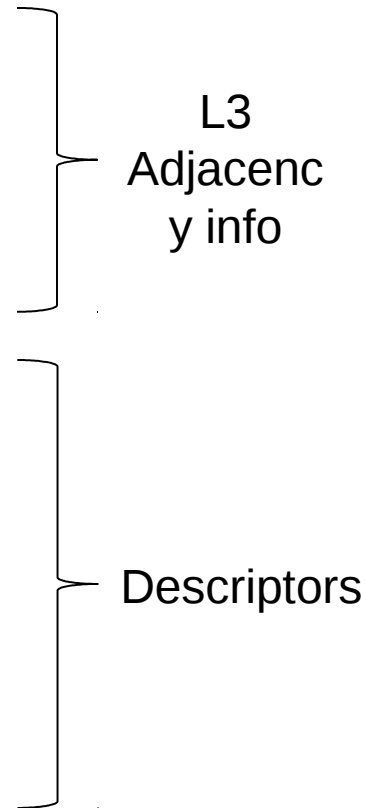
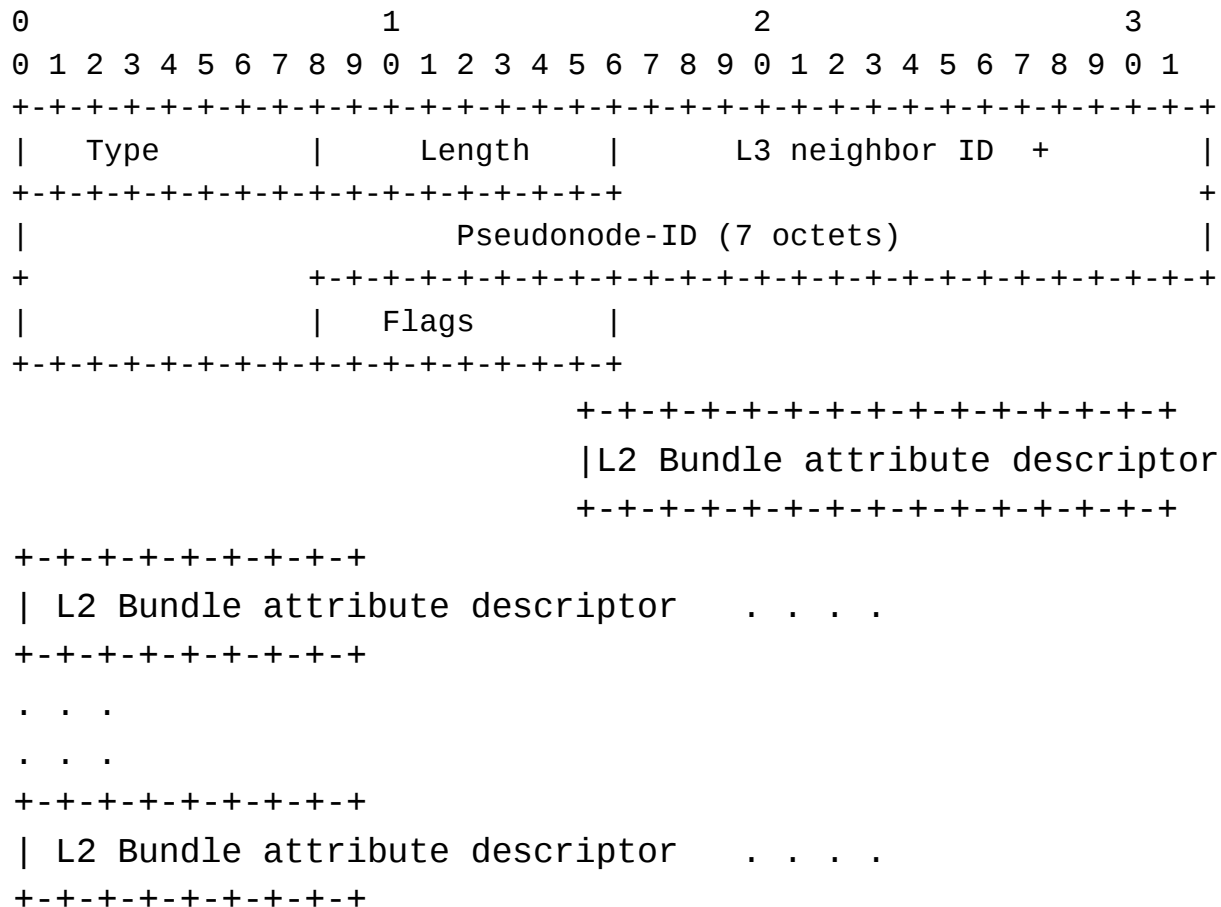
# Backup

draft-ginsberg-isis-l2bundles-00.txt

# Memory Refresh about the TLV

- Top level TLV
- Advertises individual bundle 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