

# Segment Routing BGP Egress Peer Engineering over Layer 2 Bundle

**draft-lin-idr-sr-epe-over-l2bundle-05**

Presenter: Mengxiao Chen (New H3C Technologies)  
Co-authors: Changwang Lin (New H3C Technologies)  
Zhenqiang Li (China Mobile)  
Ran Pang (China Unicom)  
Ketan Talaulikar (Cisco Systems)  
Mengxiao Chen (New H3C Technologies)

IETF-119, March 2024

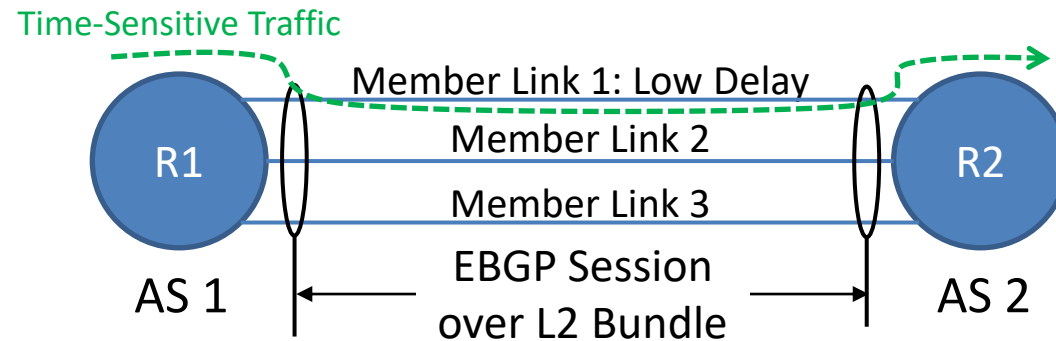
# Updates

- New co-author: Ketan.
- Main modifications after presented at IETF-118:
  - No new BGP Peering Segment Type is defined for L2 bundle member, because the definition of PeerAdj SID in RFC8402 section 4.2 does not restrict to layer 3.
  - For MPLS-SR, reuse the existing PeerAdj SID TLV for bundle members instead of defining new TLV.
  - Rewrite Section 3 "Advertising Peer Adjacency Segment for L2 Bundle Member in BGP-LS" to make it more concise.
  - Move the example to Appendix.

# Motivation

There are deployments where the Layer 3 interface on which a BGP peer session is established is a Layer 2 interface bundle (L2 Bundle).

The operator of AS1 wishes to apply a BGP-EPE policy to steer the time-sensitive traffic from AS1 to AS2 via member link 1 of the Layer 2 bundle.



BGP Peering SIDs need to be allocated to individual bundle member links, and advertisement of such BGP Peering SIDs in BGP-LS is required.

# Peer Adjacency Segment for L2 Bundle Member Link

BGP peering segments are generally advertised in BGP-LS from a BGP node along with its peering topology information, in order to enable computation of efficient BGP-EPE policies and strategies.

When a BGP peer session is established over a Layer 2 interface bundle, an implementation *MAY* allocate one or more Peer Adjacency Segments for each member link. If so, it *SHOULD* advertise the Peer Adjacency Segments of bundle members in BGP-LS, using the method defined in this document.

# Advertising in BGP-LS

- Link NLRI (parent L3 link)
- Link Attributes:
  - ① **MPLS-SR:**
    - PeerAdj SID TLV (Label for parent L3 link)
    - L2 Bundle Member Attribute TLV (member link 1) [RFC9085]
      - PeerAdj SID TLV (Label for member link 1) [RFC9086]
    - L2 Bundle Member Attribute TLV (member link 2) [RFC9085]
      - PeerAdj SID TLV (Label for member link 2) [RFC9086]
  - ② **SRv6:**
    - SRv6 End.X SID TLV (SID for parent L3 link)
    - L2 Bundle Member Attribute TLV (member link 1) [RFC9085]
      - SRv6 End.X SID TLV (SRv6 SID for member link 1) [RFC9514]
    - L2 Bundle Member Attribute TLV (member link 2) [RFC9085]
      - SRv6 End.X SID TLV (SRv6 SID for member link 2) [RFC9514]

Section 5.2 of RFC 9086

This document updates [RFC9085] and [RFC9086] to allow the PeerAdj SID TLV to be included as a sub-TLV of the L2 Bundle Member Attributes TLV.

Section 4.1 of RFC 9514

The inclusion of a L2 Bundle Member Attributes TLV implies that the identified member link is up. If any member link fails, the L2 Bundle Member Attributes TLV must be withdrawn, along with the associated Peer Adjacency SID.

# Next Steps

- Ask for WG adoption.
- Any questions or comments are Welcomed.

**Thanks**