Anycast Prefix Segments in MPLS-based SPRING

draft-psarkar-spring-mpls-anycast-segments-00

Pushpasis Sarkar psarkar@juniper.net
Hannes Gredler hannes@juniper.net

Summary

- Definitions
- Problem Statement
- Proposed Solution

Definitions

- Anycast address
  - Identifies a group devices in a given topology.
  - Packets addressed to such a anycast address from a given source device are forwarded to one or more topologically nearest node in a group of potential receiving devices
- Anycast Prefix Segments
  - Introduced by [I-D.ietf-spring-segment-routing].
  - Prefix segments associated with a anycast-address
    - Advertised by all devices participating in anycast group.

Problem

- [I-D.ietf-spring-segment-routing] introduced the use of anycast prefix segments
  - Enables usage of various use-cases along with.
  - But, when **applied to MPLS data-plane**,
    - **All devices** advertising any anycast prefix segment **MUST use/advertise the same SRGB label-ranges**.
    - **MAY NOT be possible in multi-vendor network deployments**,
      - Good mix-and-match of devices from different vendors with varying hardware/platform limitations.
    - Violates RFC 3031 as well.

Problem Statement

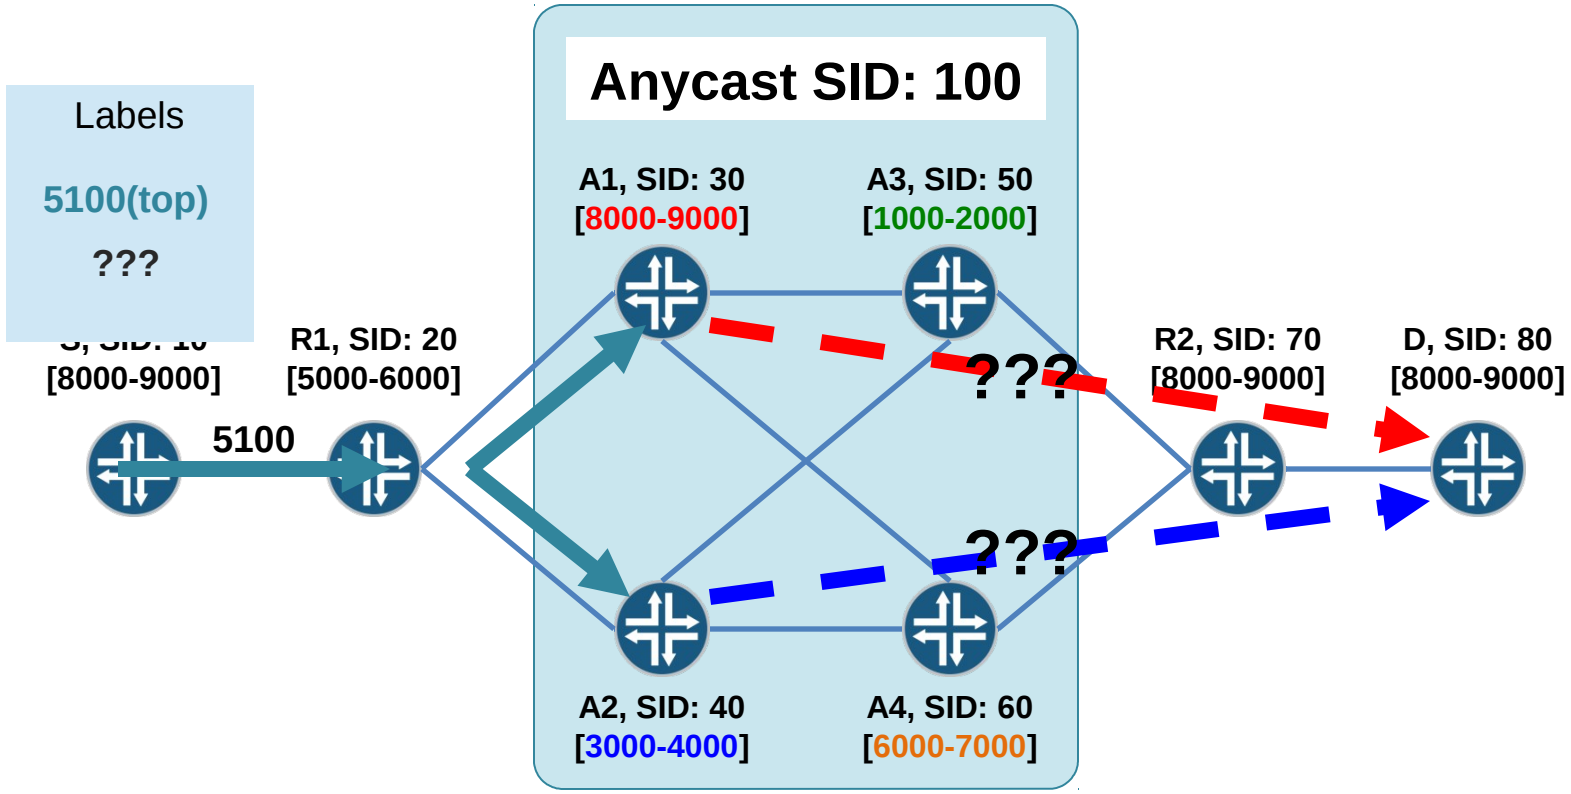

How to compute the label that represents the next segment?

Proposed Solution

- Step 1: Devices originating any anycast prefix segments
  - Create a Virtual L-FIB lookup table
    - Map all remotely learnt node/anycast prefix segment index to corresponding downstream label and next-hop.

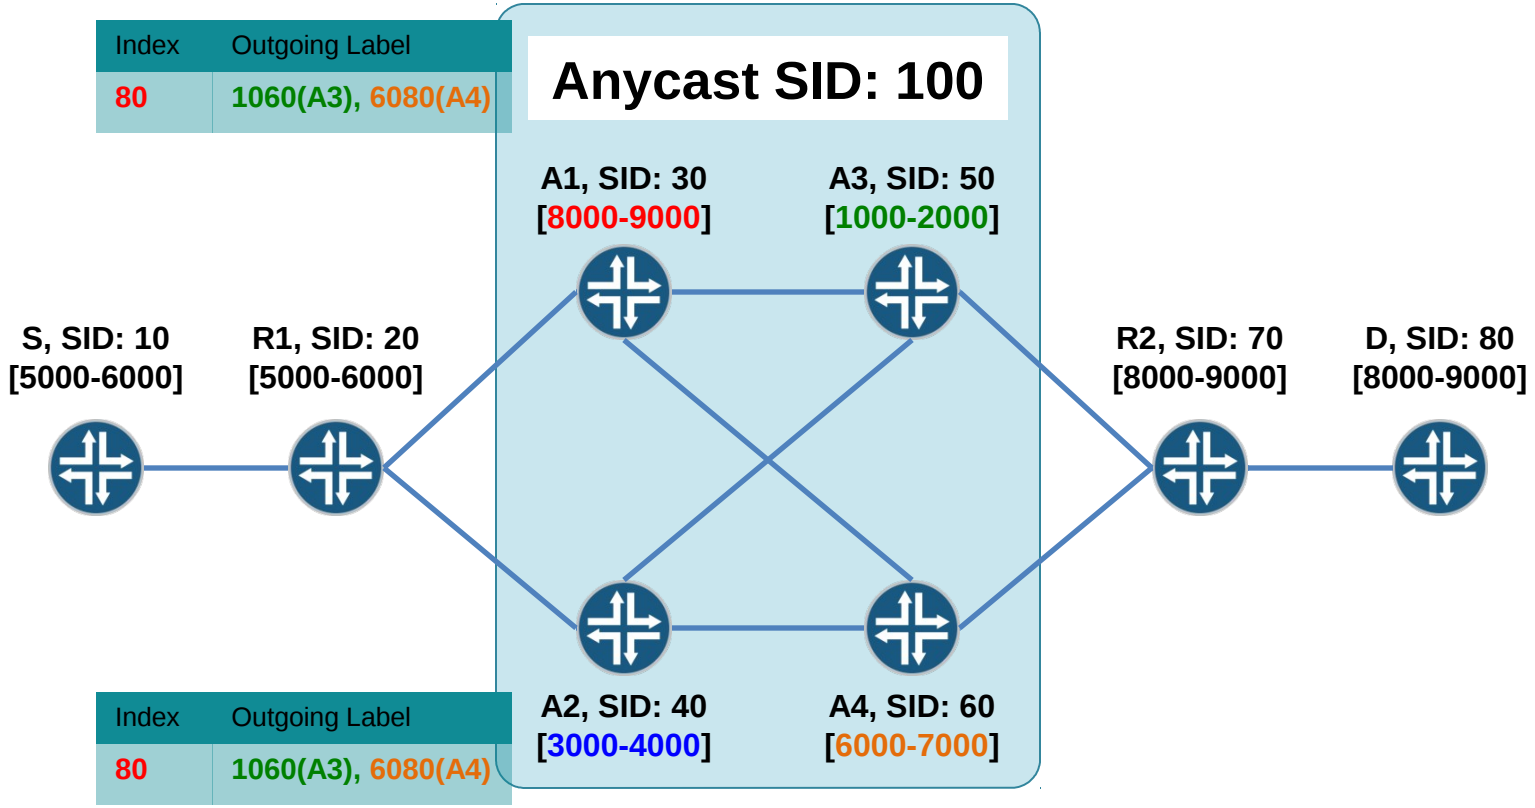

Proposed Solution

- Step 2: Devices originating any anycast prefix segments
  - **Originate IGP advertisement** for anycast prefix SID with (**No-PHP =1 and Exp-Null = 0**).
  - Ensures the packet arrives with anycast prefix segment label allocated for it. **Penultimate-hop does not POP** the label, but replaces it.

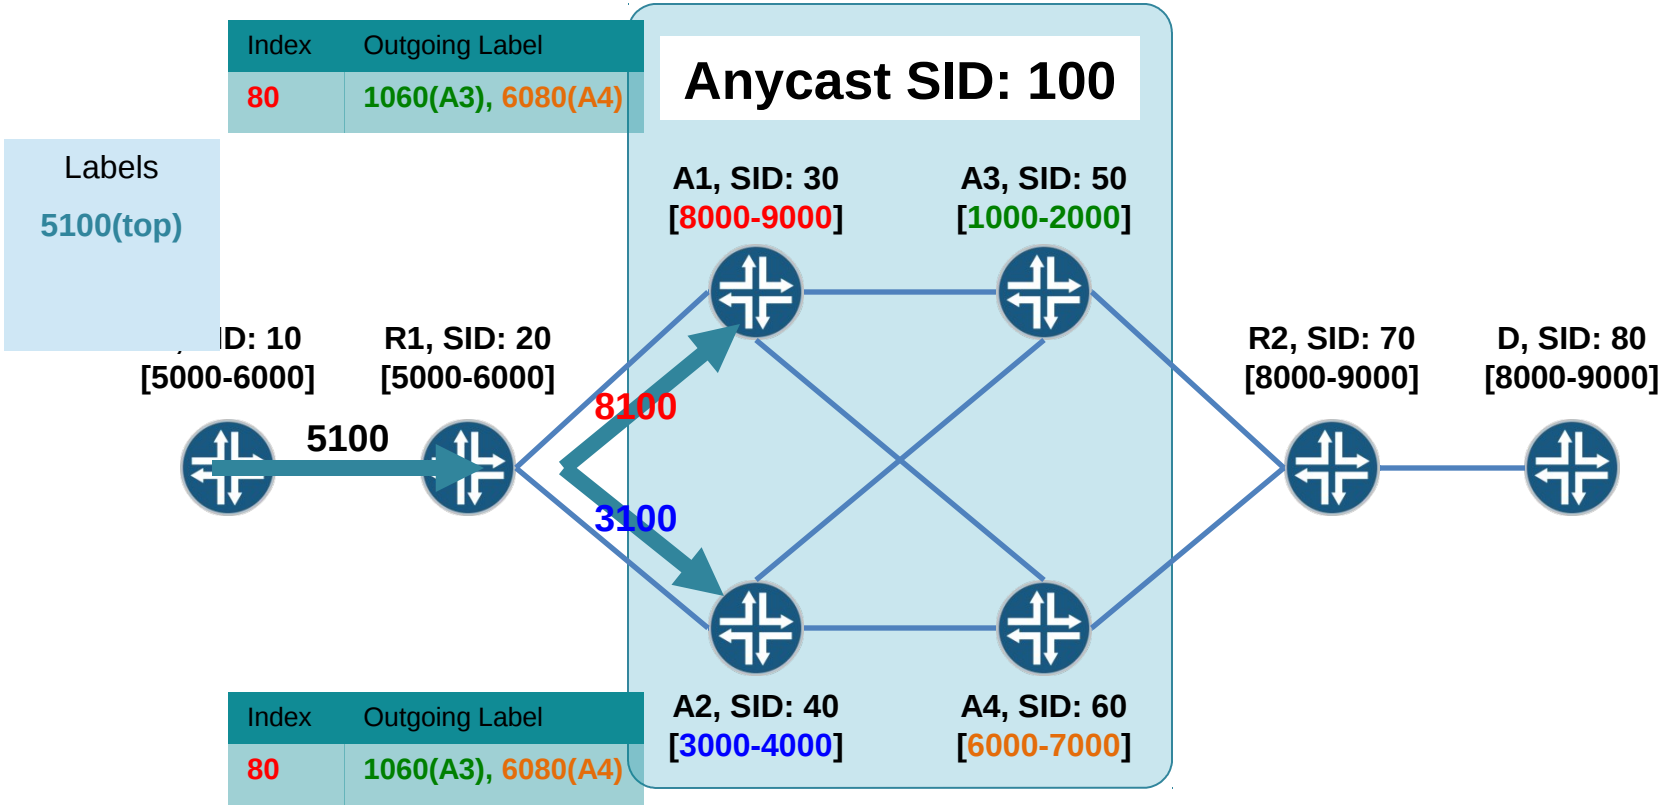

Proposed Solution

- Step 3: Devices originating any anycast prefix segments
  - **For the anycast segment label in the global LFIB table**.
    - Install a **Lookup into the Virtual LFIB** created in Step 1.

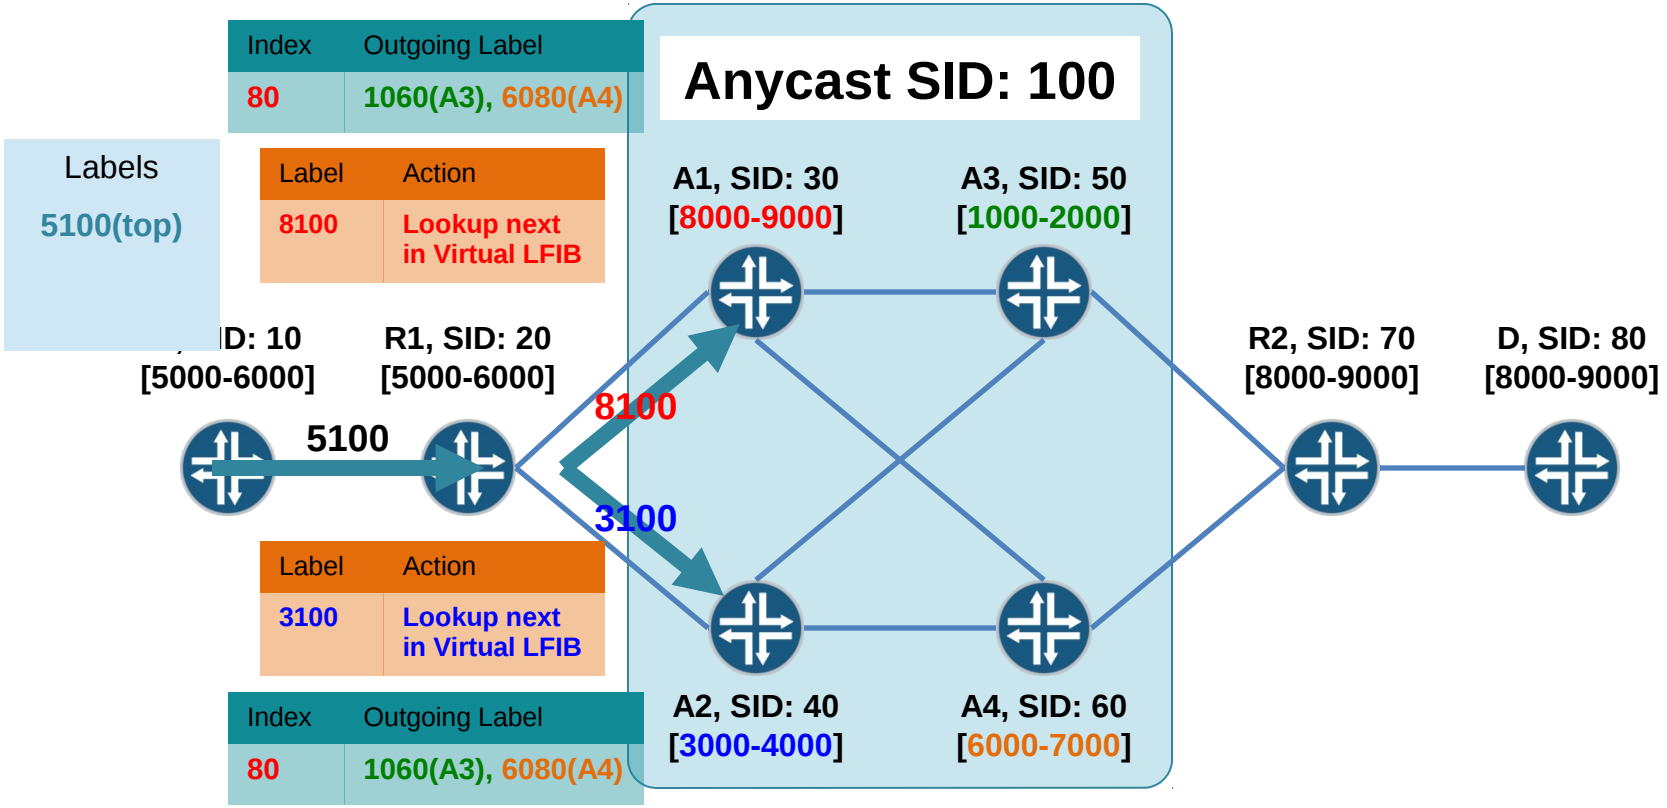

Proposed Solution

- Step 4: Ingress device using Anycast prefix segments
  - For the **prefix segment next that follows a anycast prefix segment**.
    - **Use the prefix segment index directly**. No need to compute an actual label.

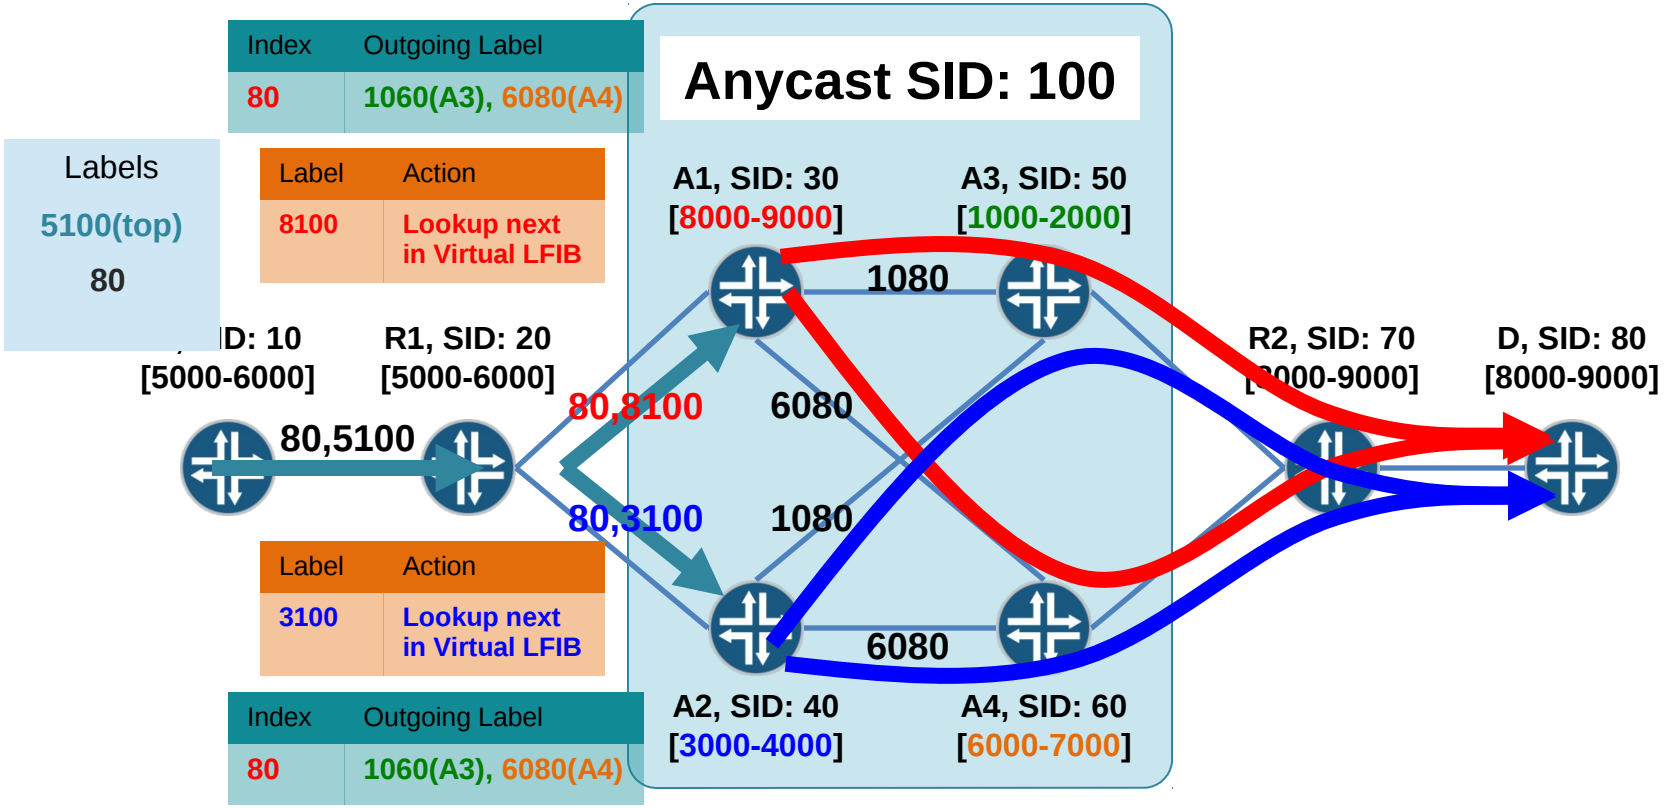

Limitations

- All devices implementing anycast prefix segments needs to **support recursive label lookup in the forwarding plane**.
  - Not a concern. Almost all chipsets do support recursive label lookup up to 3 labels.

Next Steps

- Comments/Questions/Suggestions ?

THANK YOU