P2MP Transport Using Chain Replication in Segment Routing

draft-shen-spring-p2mp-transport-chain-03

Yimin Shen, Jeffrey Zhang - Juniper Networks

Rishabh Parekh - Cisco Systems

Hooman Bidgoli - Nokia

Yuji Kamite - NTT Communications

IETF 109, Nov 2020

Motivation – Stateless P2MP Transport

Traditional P2MP in SR

Ingress replication

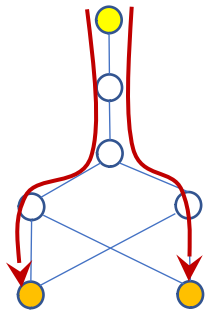

- P2P tunnels, one per leaf
- Stateless in core
- No traffic optimization

Controller-driven P2MP tree

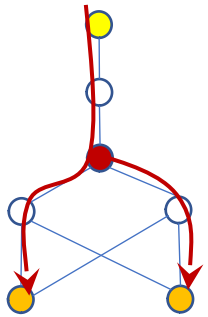

- Max traffic optimization
- Controller-provisioned tree state on branch nodes
- Stateful in core

Stateless P2MP

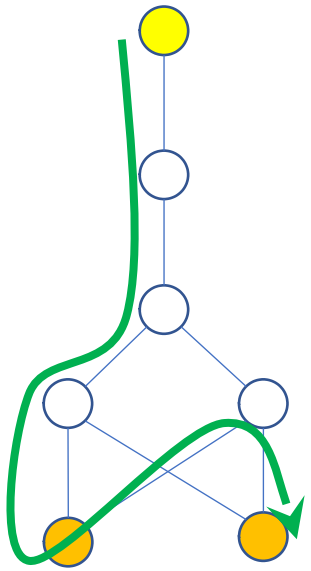

- P2MP chain tunnels
- ✓ One P2MP chain reaches multiple leaf nodes
- ✓ Provisioning on root
- ✓ Traffic optimization
- ✓ Stateless in core

P2MP Chain Tunnels

A P2MP chain is a single-path tunnel that reaches multiple leaf nodes.

Root sends packets over one or a small number of P2MP chains.

Applicable to all topologies. Most beneficial for:

1) Ring topology: A single P2MP chain per multicast stream.
2) Linear topology: Reduced traffic and tunnels across domains.

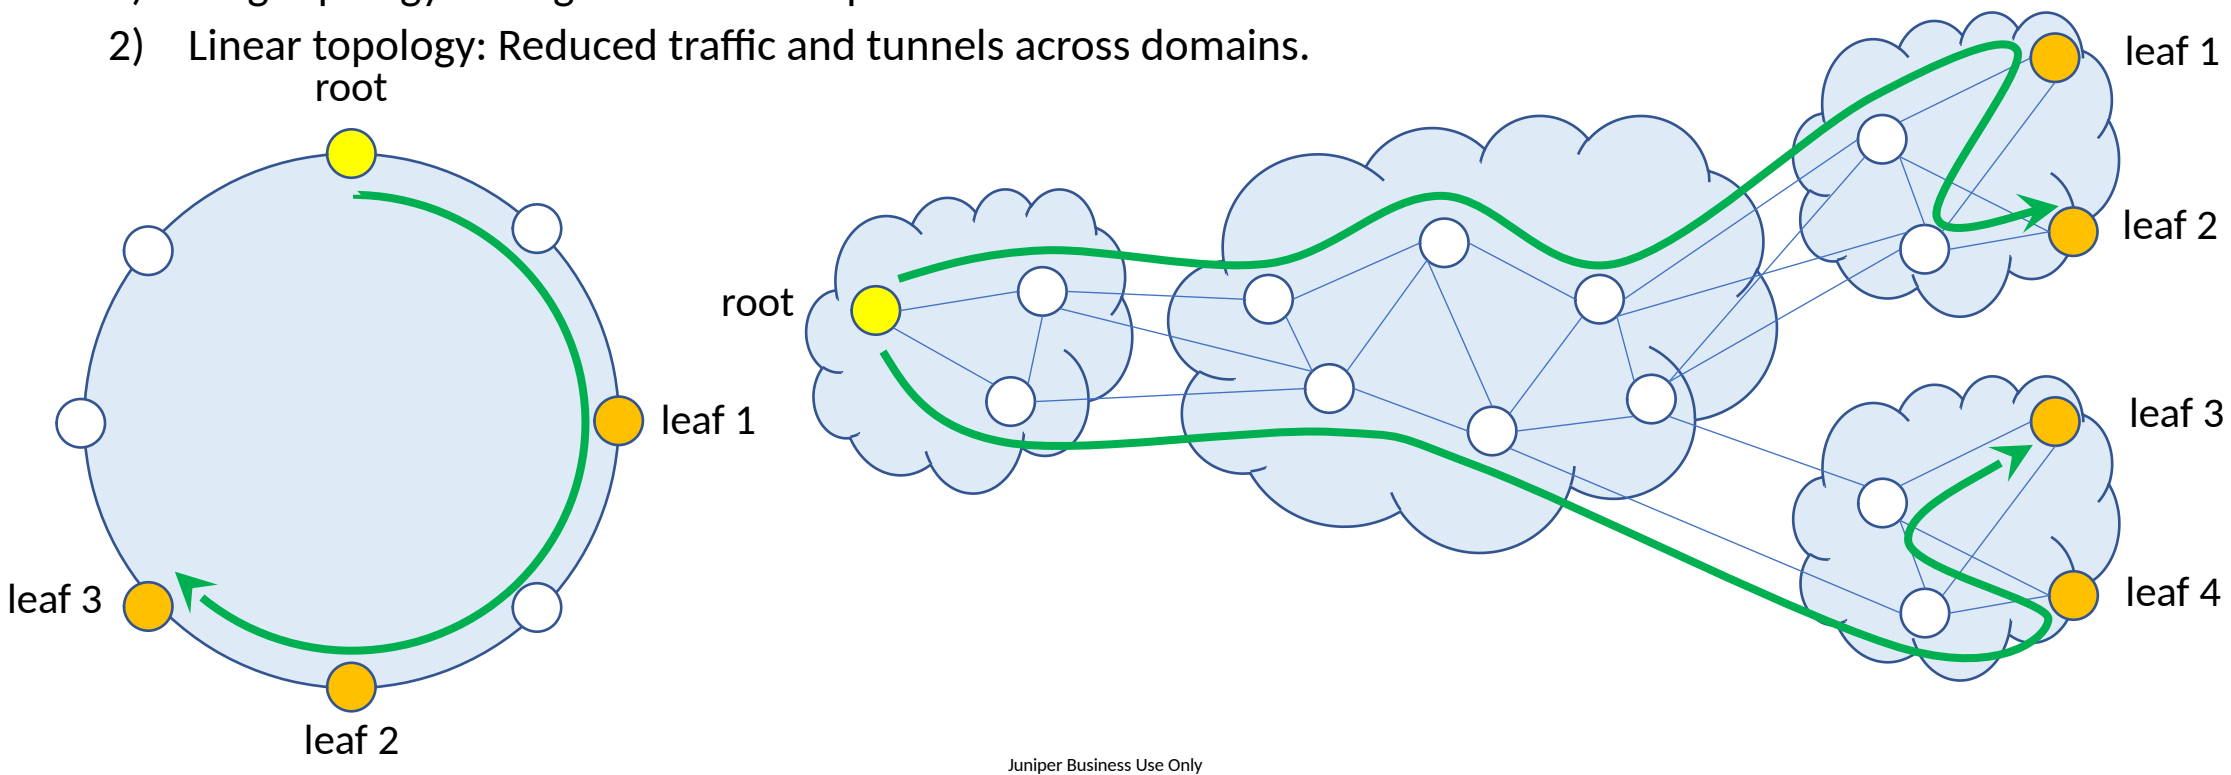

P2MP Chain and Chain Replication

One tail-end leaf node is a normal receiver.

One or multiple transit leaf nodes, aka. **bud nodes**.

- Chain replication + local processing + forwarding
- Modeled as **bud segments**, with **bud-SIDs**.

A P2MP chain comprises a SID list with bud nodes represented by bud-SIDs.

- Provisioned on root node in the same manner as a point-to-point tunnel.

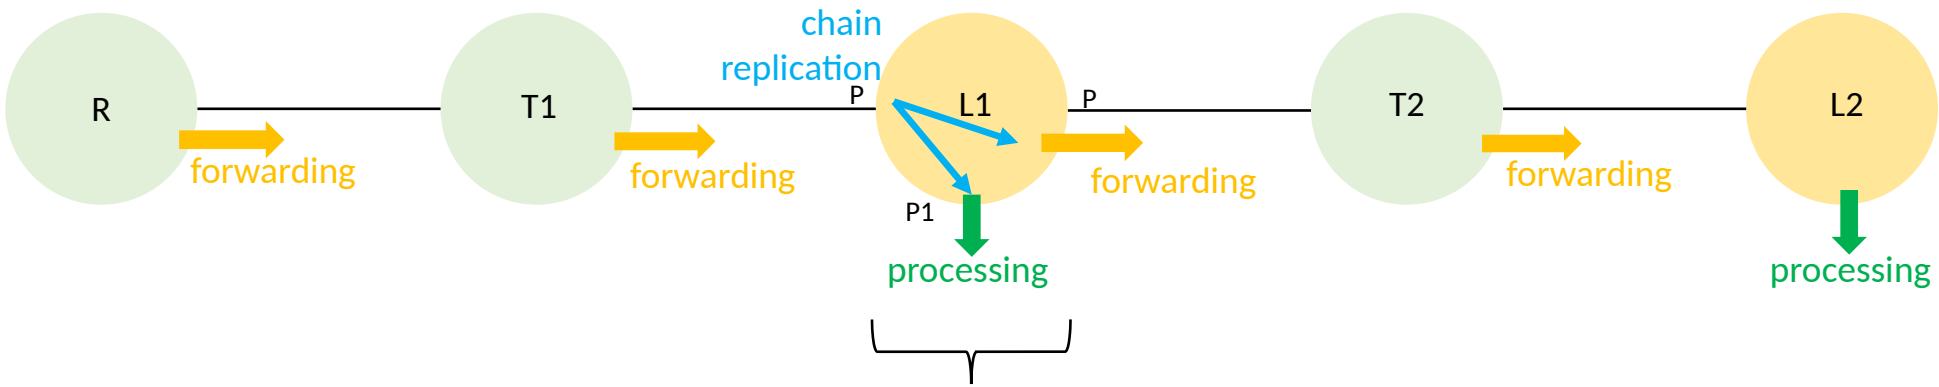

bud segment : replicate P → P1
forward P down the chain
process P1 locally

Bud Segments

Nodal segments on each router

- One for SR-MPLS; One for SRv6

Global segments

- Bud-SIDs are allocated from SRGB

Routable segments via the shortest paths

- Can be used along with other segments to build explicit paths

Can be advertised by ISIS, OSPF, and BGP

Sharable building blocks for stateless P2MP tunnels

Example 1

A1 – A7: adj-SIDs

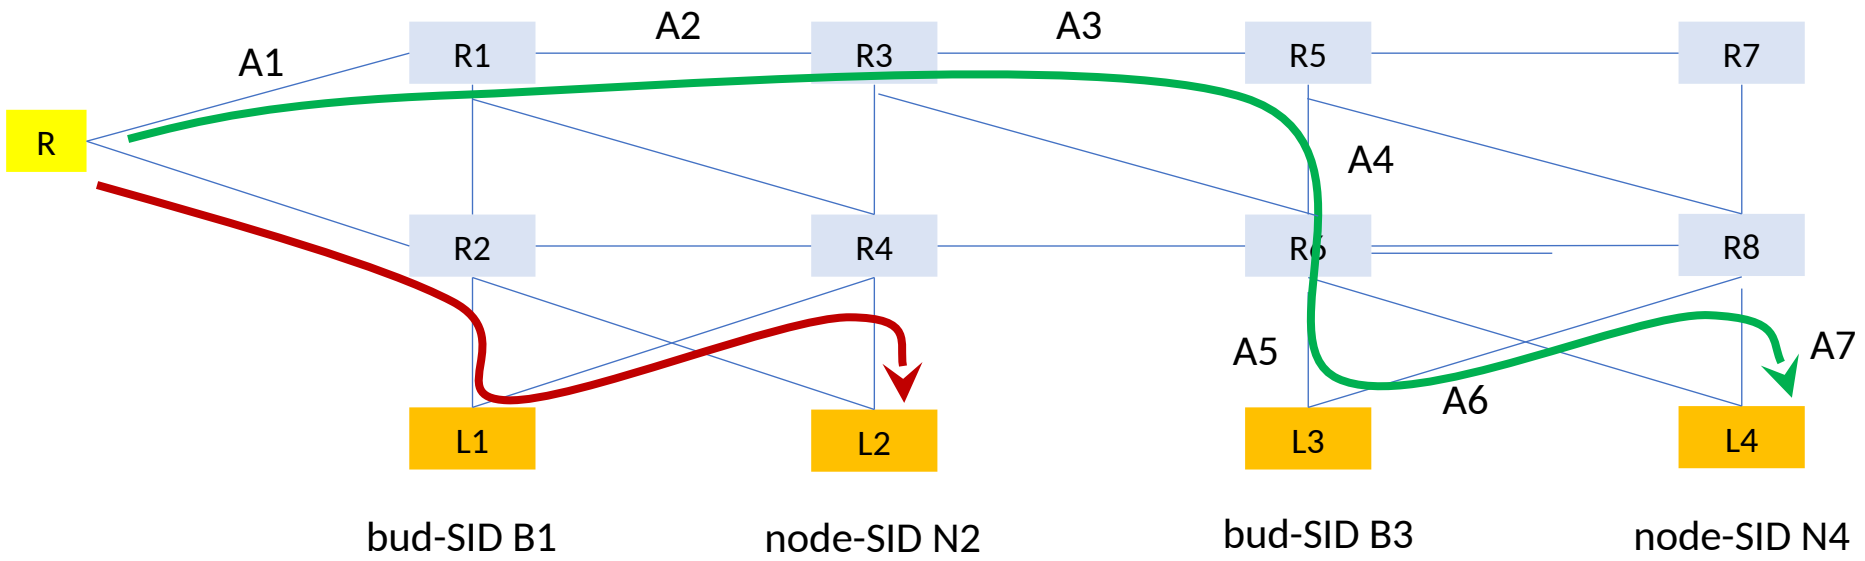

A multicast stream to L1, L2, L3, and L4, using two P2MP chains:

- Red chain to L1 and L2 takes the shortest path from R to L1, and from L1 to L2.
  - SID list = {B1, N2}.
- Green chain to L3 and L4 takes an explicit path from R to L3, and from L3 to L4.
  - SID list = {A1, A2, A3, A4, A5, B3, A6, A7}

Example 2

Green stream

- root = R0, leaves = R1, R3, R5, R7
- SID list = {B_1, B_3, B_5, N_7}

Red stream

- root = R1, leaves = R2, R5, R6
- SID list = {B_2, B_5, N_6 }

R5 is a leaf node of both streams.

- B_5 appears in both P2MP chains.

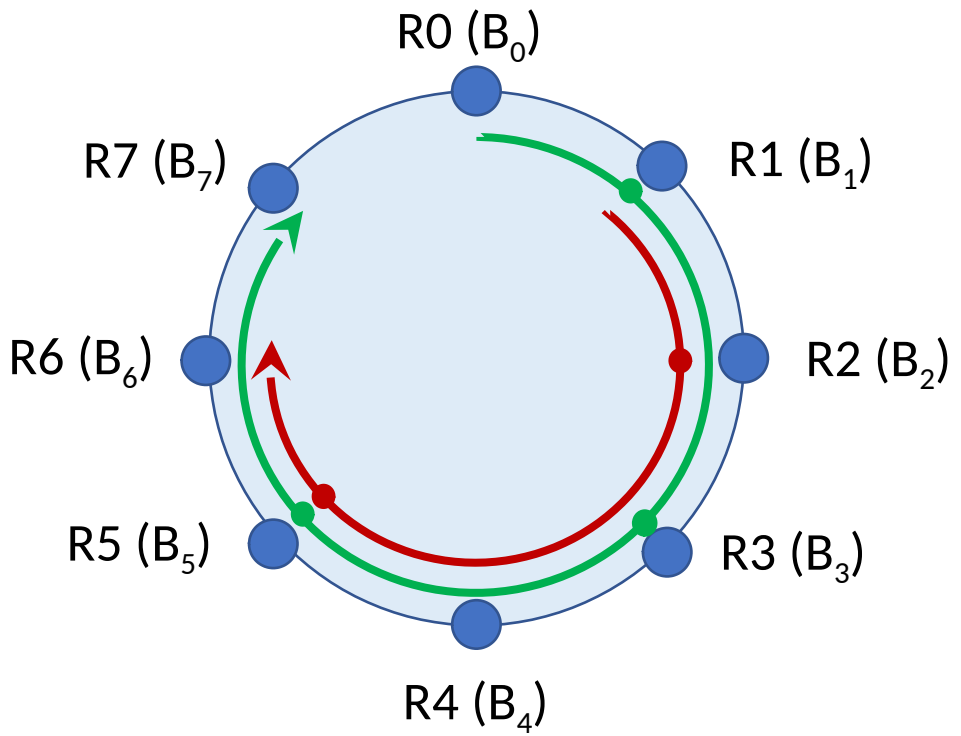

General Model of Bud Segment Behavior

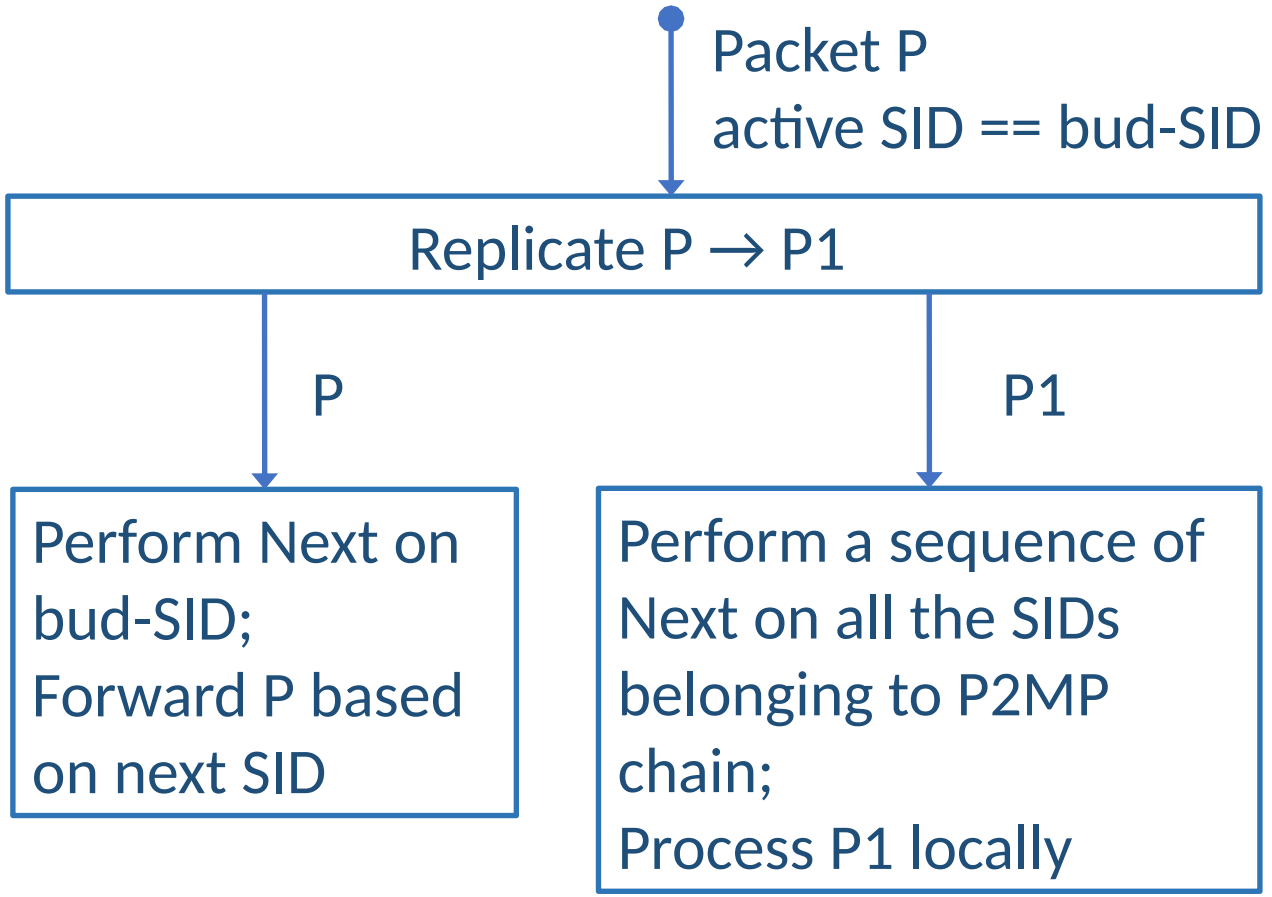

Bud Segment Behavior in SR-MPLS

MPLS header: P2MP chain labels + EoC + service label (optional; co-existing)

- If there is a service label, root inserts an End-of-Chain (EoC) label after P2MP chain labels.
- EoC is a new Extended Special-Purpose Label, i.e. <XL = 15, EoC>.

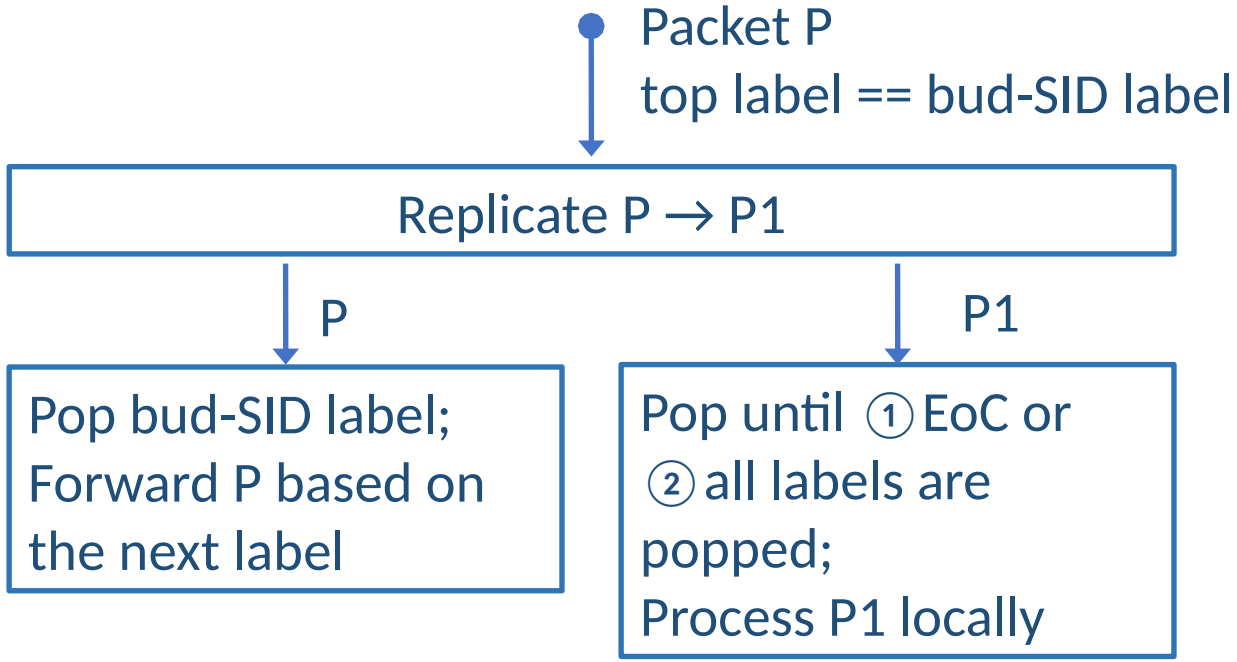

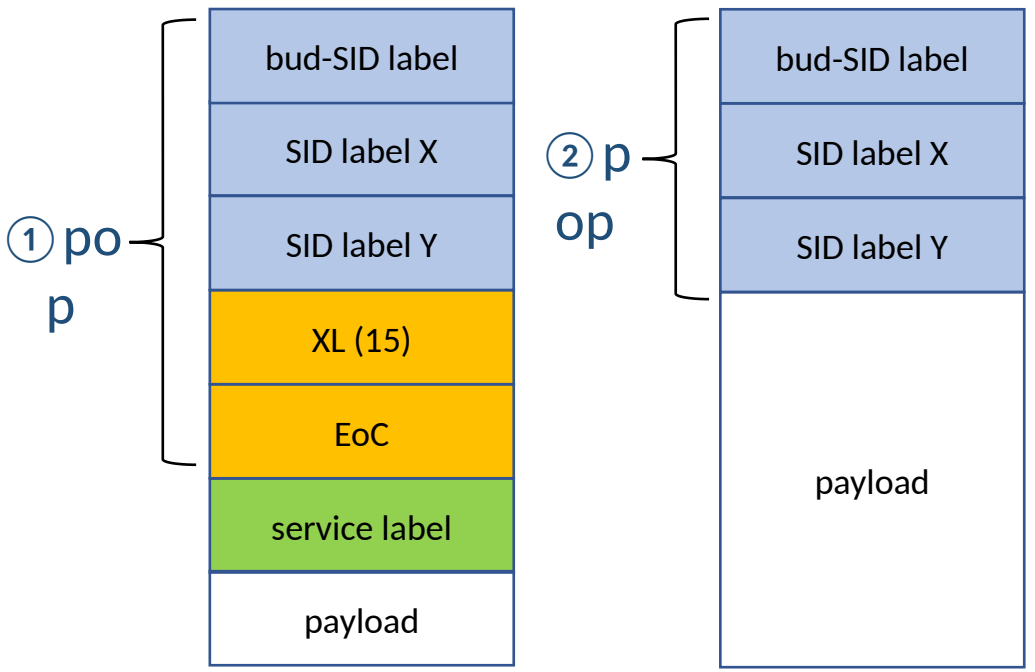

Example: EoC Label and Service Label

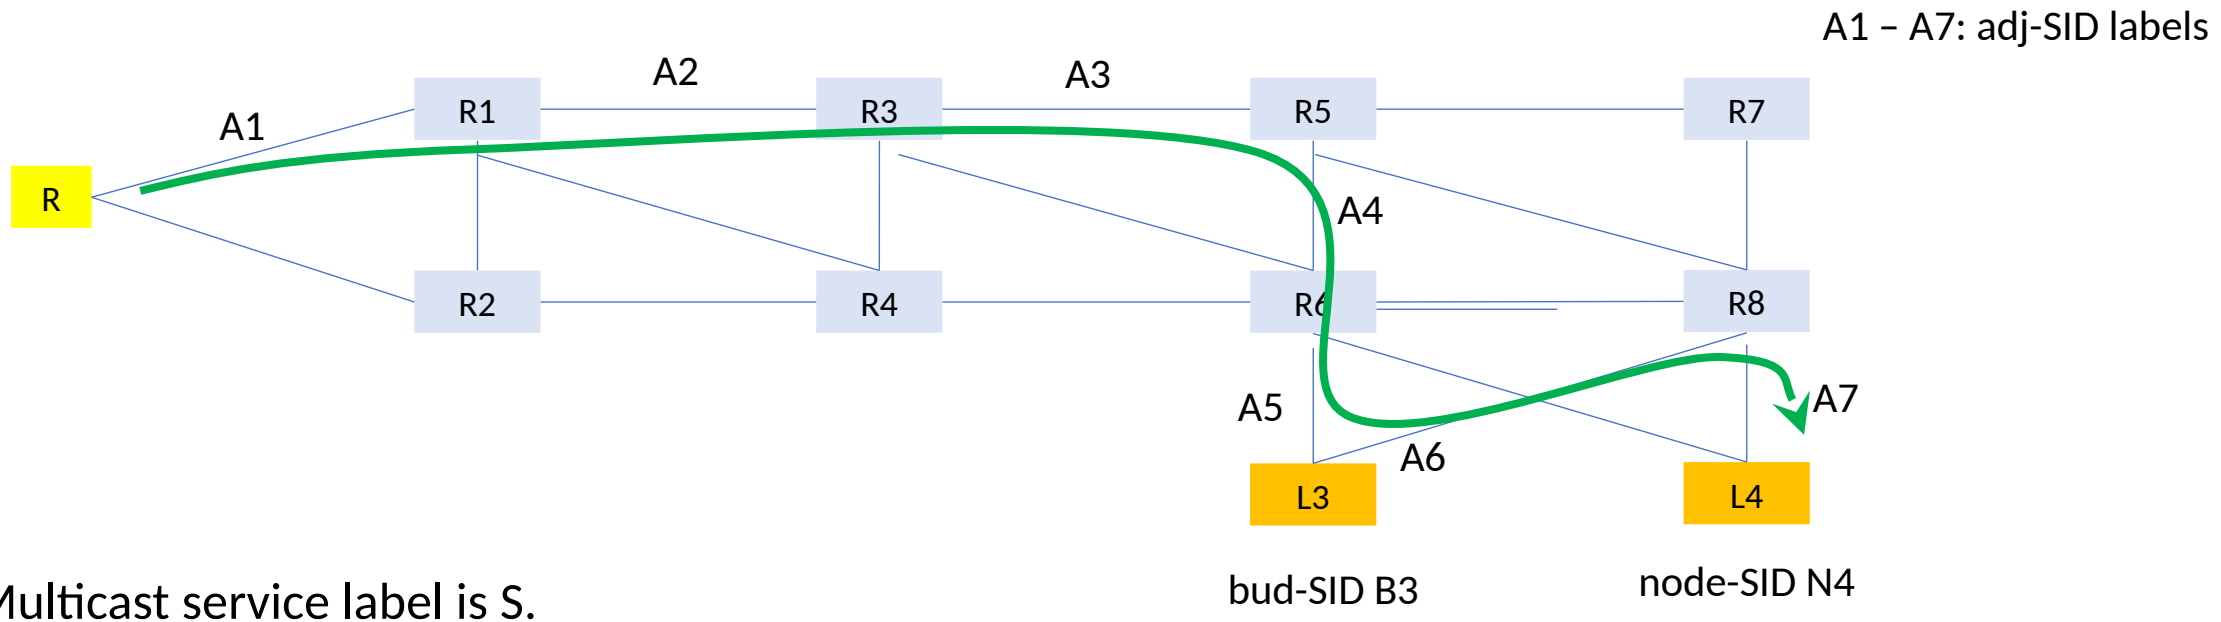

Multicast service label is S.

P2MP chain's SID list: {A1, A2, A3, A4, A5, B3, N4} .

R sends a service packet P with {A2, A3, A4, A5, B3, N4, XL, EoC, S}.

L3 receives P with {B3, N4, XL, EoC, S}, and replicates it to generate P1.

- For P, pops B3, and forwards the packet with {N4, XL, EoC, S}.
- For P1, pops B3, N4, XL, and EoC, and processes the packet with {S}.

L4 receives P with {N4, XL, EoC, S}, pops N4, XL, and EoC, and processes the packet with {s}.

Bud Segment Behavior in SRv6

IPv6 header + SRH (P2MP chain SIDs) + IP/L2 header + payload

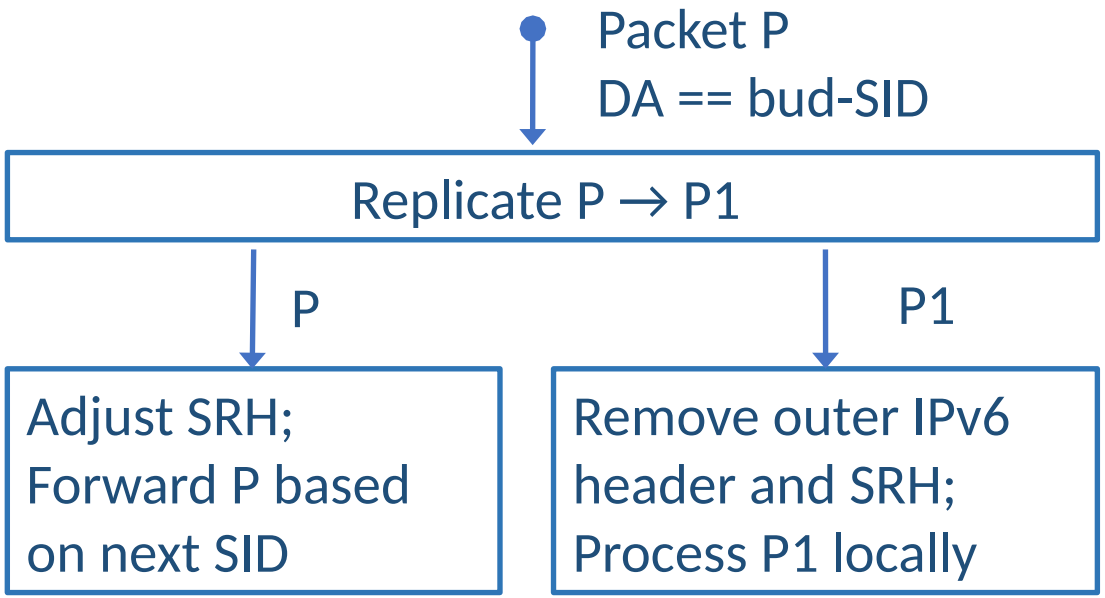

Path Computation

P2MP chain computation is single-path computation

- Can use algorithms extended from P2P path computation.

Specific considerations

- Max hops – limit total delay
- Max hops between two consecutive leaf nodes – avoid sparse chain
- Max times that a link or node may be traversed – avoid efficiency degradation
- Leaf groups based on location or policy.
  - A group is a sequence or set of leaf nodes, treated as loose hops.
  - A P2MP chain is computed for each leaf group.

Next Steps

- Welcome feedback from the WG
- Specify ISIS/OSPF/BGP extensions for bud segment advertisement

Thank you