SRv6 Segment List Optimization

draft-liu-idr-srv6-segment-list-optimize-00

Presenter: Changwang Lin

Co-authors: Yisong Liu(China Mobile)
Changwang Lin (New H3C Technologies)
Ran Chen (ZTE Corporation)
Yuanxiang Qiu(New H3C Technologies)

IETF-117

Background

In the following scenarios, it is necessary to specify the End SID of the egress node in the Segment List of the SRv6 Policy.

- Steering traffic based on binding SID

  For example, in tunnel splicing scenarios and cross-domain path splicing scenarios, requires specifying the End SID of the egress node in

  the Segment List of the SRv6 Policy.

- End-to-End fast fault detection based on SRv6 policy

  The END SID of the egress node must be specified in the Segment List of OAM messages.

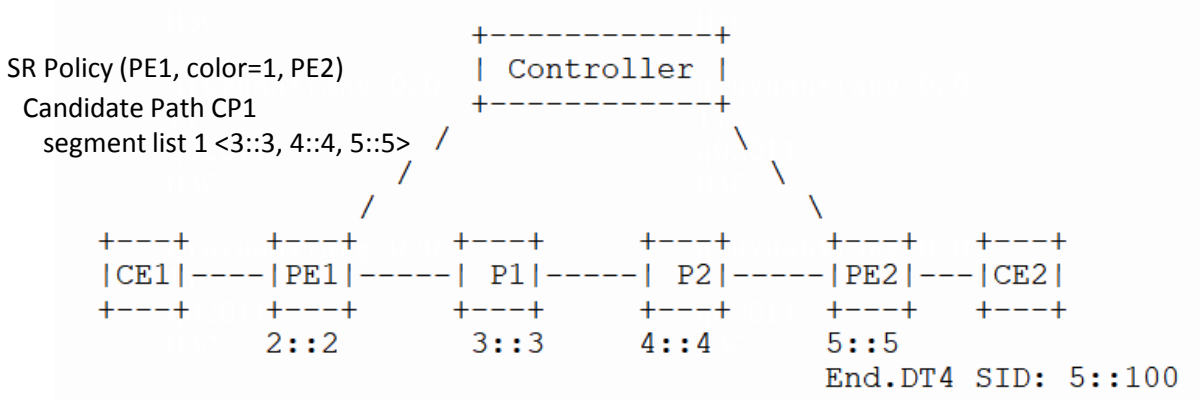

SRH of packets from PE1 to PE2:

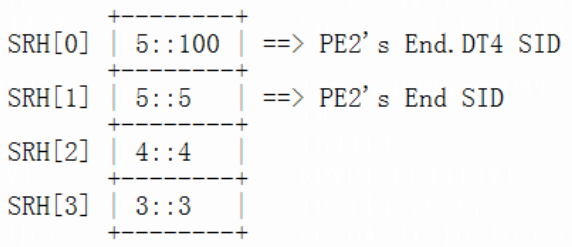

Data packet from PE1 to PE2

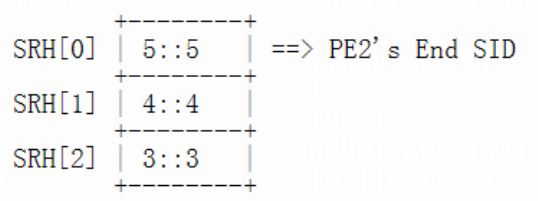

OAM packet of SR Policy(PE1, 1, PE2)

The SRH (Segment Routing Header) of data packets forwarded based on the SRv6 Policy will encapsulate both the End SID and VPN Service SID of the egress node.

Issues

This encapsulation will result in the following two issues:

- PSP behavior may not be feasible.

  PSP behavior may not be feasible since the condition of (SL==0) is not met, which means the penultimate SR Segment Endpoint Node cannot remove the SRH from the IPv6 extension header during processing.

- The forwarding efficiency of the egress node decreases

  The data packet will need to look up the SID (Segment Identifiers) table twice.For certain chips, implementing the second SID table lookup requires a loopback interface. Due to bandwidth limitations and the possibility of other service packets being in the loopback interface, packet forwarding efficiency towards VPNs will be greatly affected.

Proposal

When both the End SID and VPN Service SID of Egress Node are present,
the End SID can be removed while forwarding packets.

- The controller disseminates a segment list via BGP and indicates which SID belongs to the egress node. It notifies the ingress node by extending the Segment Flags of the Segment Types sub-TLVs.

- The ingress node optimizes the SRH.SegmentList of packets.
  - When there are End SID and service SID of egress node on the path,
  - and if SRH.SegmentList of the packet already contains the service SID of the egress node,

  the End SID of the egress node will not be encapsulated in the segment list at the same time.

Extend the Segment Flags of Segment Types sub-TLV

- Define a bit to identify which is the egress node's SID

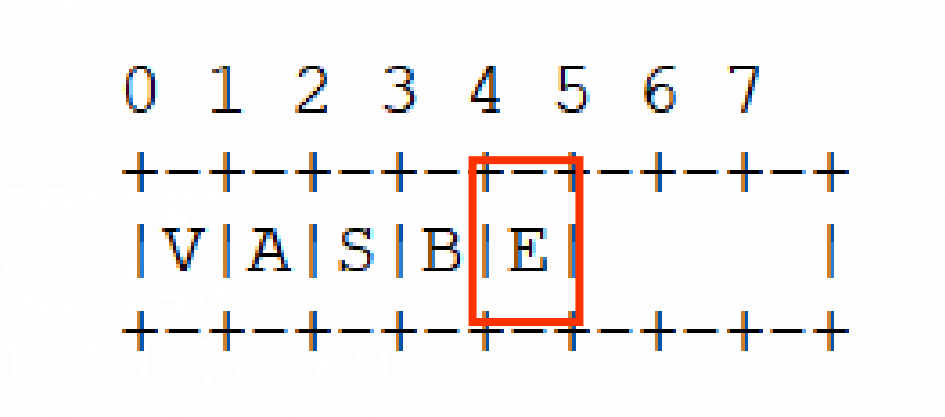

E-Flag: When set, indicates that this segment ID is the egress node's SID.

Example

SR Policy (PE1, color=1, PE2)
Candidate Path CP1
segment list 1 <3::3, 4::4, 5::5>
END SID 5::5 with E flag

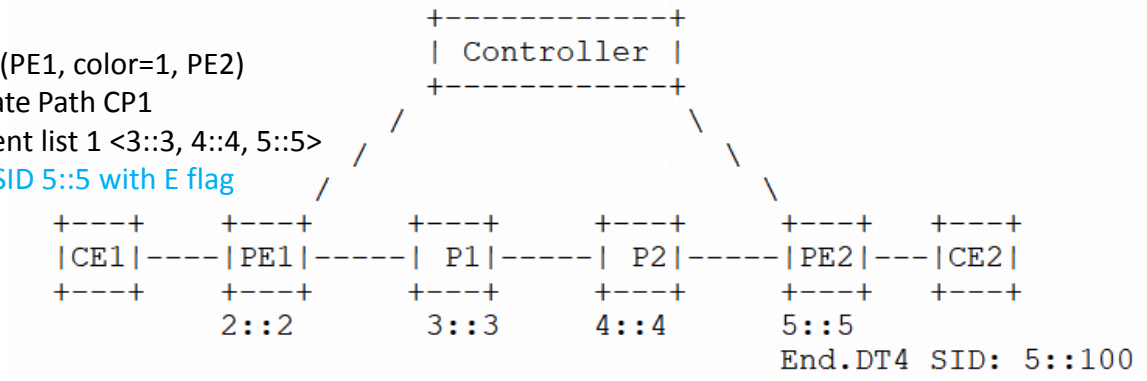

- For data packets forwarded to VPN

  SRH will not encapsulate the End SID corresponding to the egress node in the SID list of SRv6 Policy.

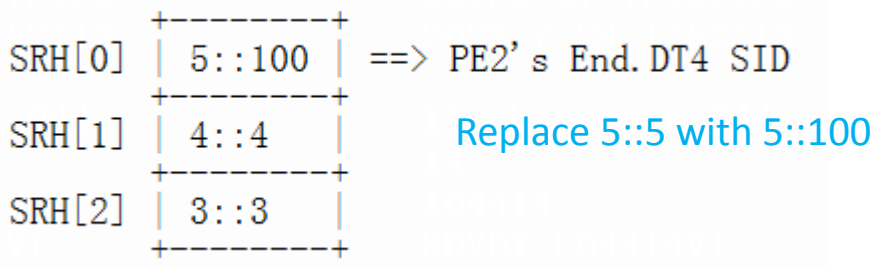

- For OAM packets

  all node SIDs of the SID lists of SR policy will be encapsulated.

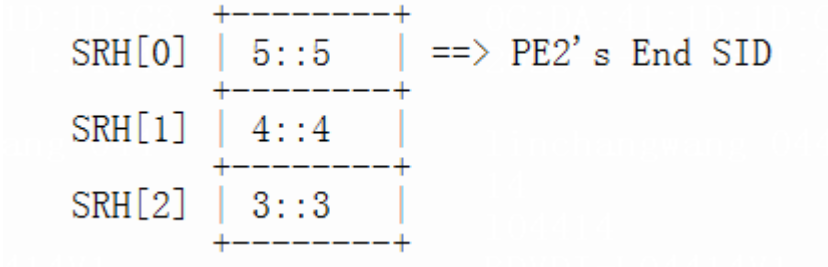

Next Steps

- Any questions or comments are Welcomed
- Seeking for feedback