

PCEP Extension to Support SRv6 Segment List Optimization

draft-lin-pce-srv6-segment-list-optimize-00

Presenter: Ran Chen

Co-authors: Changwang Lin (New H3C Technologies)

Yisong Liu(China Mobile)

Ran Chen (ZTE Corporation)

Yuanxiang Qiu(New H3C Technologies)

IETF-117

Background

In the following scenarios, the End SID of the egress node must be specified in the Segment List of SRv6 Policy.

- **Steering traffic based on binding SID**

For example, In tunnel splicing scenarios and cross domain path splicing scenarios.

- **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.

The SRH of data packets forwarded based on the SRv6 Policy will simultaneously encapsulate the End SID and VPN SID of the egress node.

Introduce the following questions:

1. PSP behavior may not be executable.

Because the condition of (SL==0) is not met, the penultimate SR Segment Endpoint Node will not be able to perform the processing of removing the SRH from the IPv6 extension header.

2. The forwarding efficiency of egress node decreases.

The data packet needs to look up the SID table twice within the egress node. For some chips, the second SID table lookup requires a loopback interface to be implemented.

Due to the bandwidth limitations and the possibility of other service packets coexisting on the loopback interface, the forwarding efficiency of packets to VPN will be greatly affected.

Proposal

- **When the controller distributes a segment list through PCEP, it informs which one is the egress node's SID.**

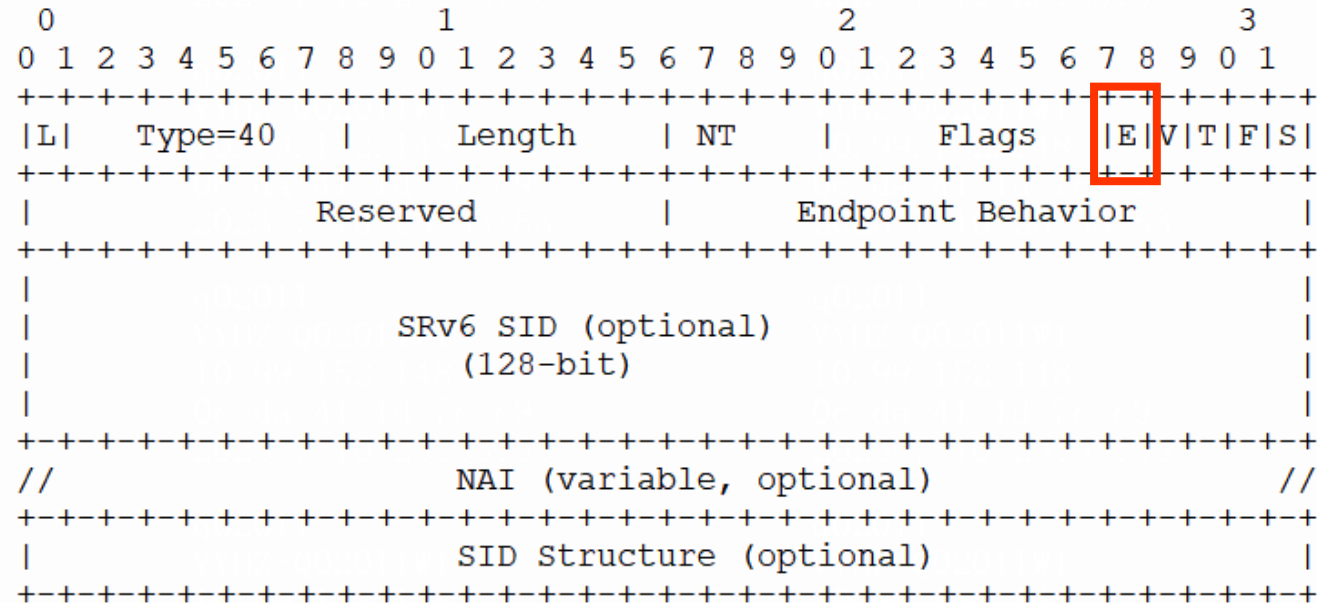
Inform the ingress node through a flag of the SRv6-ERO subobject.

- **The ingress node optimize 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.

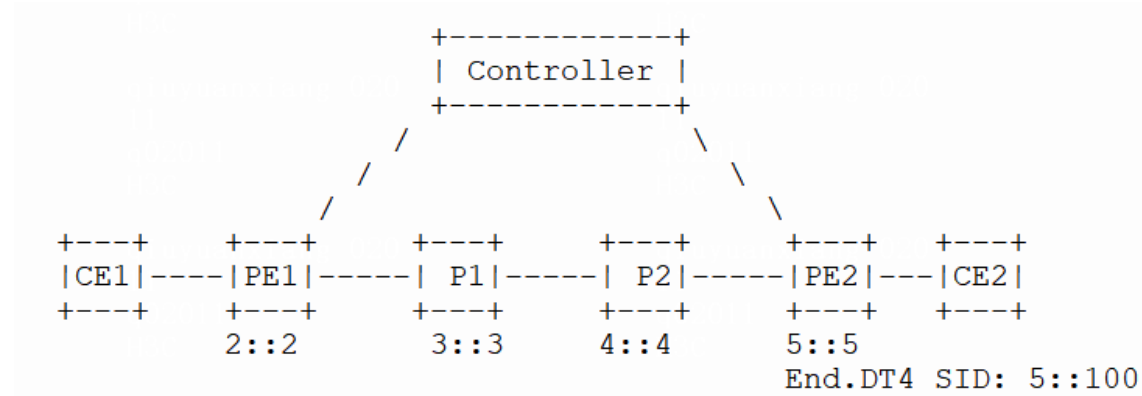
Extend the Flags of SRv6-ERO subobject

- Define a E 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



- **For data packets forwarded to VPN**

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

```

Segment List[0] | 5::100 | ==> PE2's End.DT4 SID
Segment List[1] | 4::4   |
Segment List[2] | 3::3   |
  
```

- **For OAM packets**

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

```

Segment List[0] | 5::5   | ==> PE2's End SID
Segment List[1] | 4::4   |
Segment List[2] | 3::3   |
  
```

Next Steps

- Any questions or comments are Welcomed
- Seeking for feedback

PCEP extension to support CPs validity

draft-chen-pce-sr-policy-cp-validity-00

Presenter: Ran Chen

Co-author: Ran Chen (ZTE)
Detao Zhao (ZTE)
Mike Koldychev (Cisco)


PCE WG IETF-117 Meeting, July 2023

Introduction

- SR Policy architecture are specified in [RFC9256]. An SR Policy comprises one or more candidate paths (CP) of which at a given time one and only one may be active . Each CP in turn may have one or more SID-List of which one or more may be active. when multiple SID-List are active then traffic is load balanced over them. However, a CP is valid when at least one SID-List is active.
- This candidate path validity criterion cannot meet the needs of some scenarios.
- draft-chen-spring-sr-policy-cp-validity defines the validity control parameters under candidate Path to control the validity judgment of candidate
- This document defines PCEP extensions for signaling the validity control parameters of a candidate path for an SR Policy.

Motivation

- The candidate path validity criterion defined in [RFC9256] can't meet the needs of the following scenarios:



- The CP1 carries a total of 200MB of traffic. Within the POL1, the flow-based hashing over its each SL with a ratio 50%, that is each SL carry 100MB of traffic. At this time, if one of the Segment Lists is invalids, the remaining Segment List cannot carry 200MB of traffic. However, the CP1 is still active.

Extensions

- draft-ietf-pce-segment-routing-policy-cp defines SR Policy Association (SRPA), and the SR Candidate Paths are the members of this Association.
- We define the CP validity TLV in the SR Policy Association (SRPA) object to signal the validity control parameters of a candidate path.

Next Step

- Comments welcome.

Thanks!