

BGP SR Policy Extensions for Administrative Flags

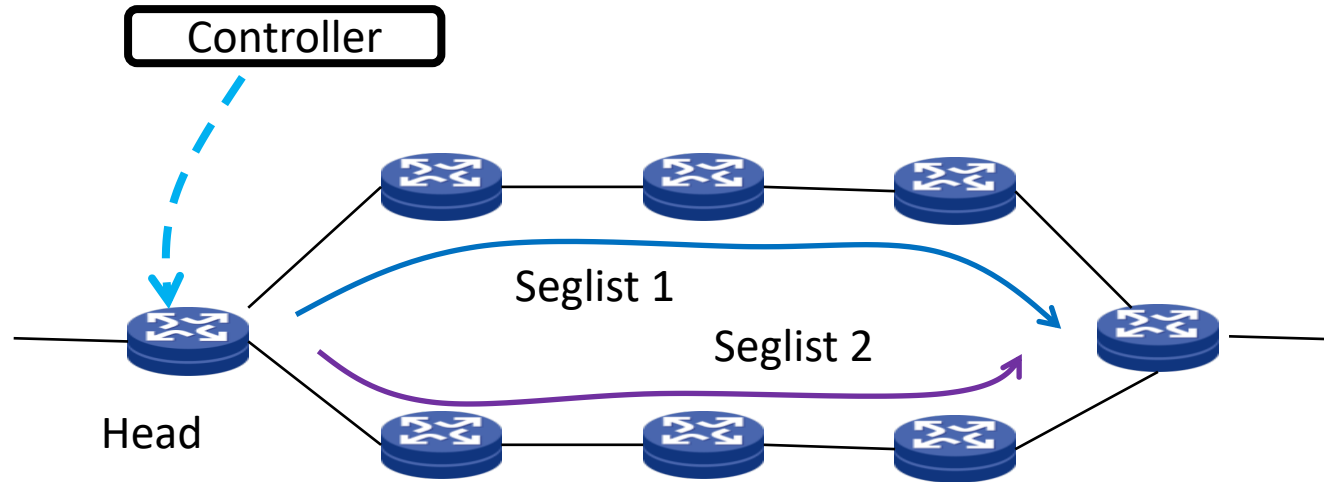
draft-lin-idr-sr-policy-admin-flags-00

Presenter : Changwang Lin (New H3C Technologies)

Co-Authors : Changwang Lin (New H3C Technologies)
Jinming Li(China Mobile)

IETF-121

Scenarios 1: shutdown & backup ineligible



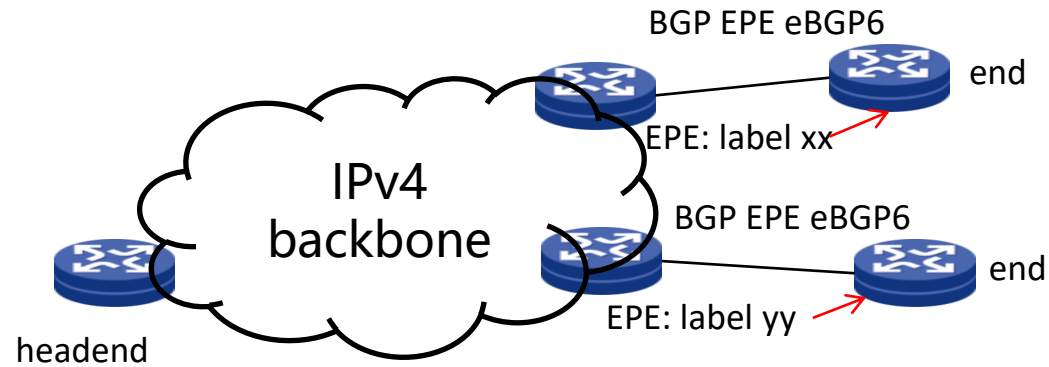
Problem:

- The controller issues an SR Policy with two Segment Lists. It just temporarily wants to shut down a path or a segment list for management and maintenance.
- In some scenarios, CP cannot be used as a backup path.

Solution:

- Using **S-Flag** in Admin Flags to **shutdown SR Policy Path or Segment List**.
- Using **B-Flag** in Admin Flags to control this CP **cannot be used as a backup path**.

Scenarios 2: Ignore service route's Prefix SID




Problem:

- In the **6EPE** scenario, the SR-Policy path **already includes the EPE label**.
- In encapsulating at the headend, even though the segment list of the SR policy can already reach the egress node, the service SID for the egress node is still added, resulting in an additional service SID.

Solution:

- Using **I-Flag** : Indicates the CP is configured as "Ignore service route's Prefix SID". It allows traffic to a BGP service route to be steered over an SR policy without imposing the service route's prefix label or SRv6 Service SID.

Scenarios 3 Drop-Upon-Invalid behavior



Problem:

- In RFC 9256, section 8.2 defines the Drop-Upon-Invalid behavior. An SR Policy MAY be enabled for the Drop-Upon-Invalid behavior. Currently, there is no behavior control for Drop-Upon-Invalid on the path of an SR policy. In draft-ietf-idr-sr-policy-safi-09, the Binding SID Flags, this flag is not suitable for controlling path behavior.

Solution:

- The D flag encodes the "Drop Upon Invalid" behavior and is used by a path of an SR policy, as described in section 8.2 of [RFC9256].

Scenarios 4 SR Segment List optimization

Ref to [draft-liu-idr-sr-segment-list-optimize-00]: [Discuss whether to define a CP or a Segment List flag.](#)

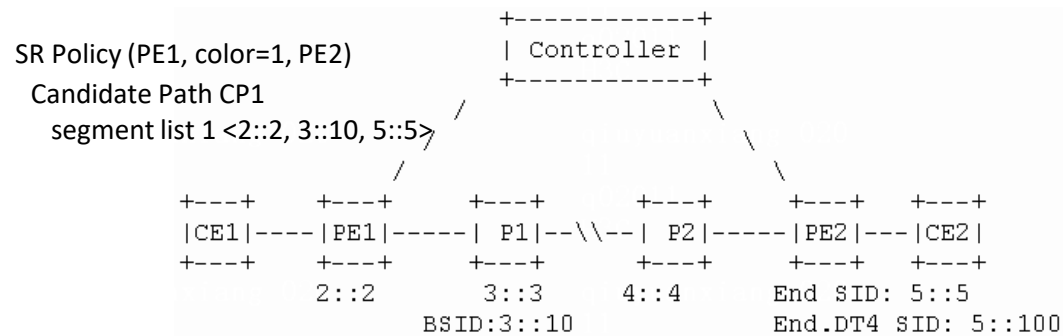
In the following scenarios, it is necessary to specify whether the last SID in the Segment List of the SRv6 Policy is the End SID of the egress node, or whether the candidate path includes or excludes the last SID.

- **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.SegmentList of packets from PE1 to PE2:



Data packet from CE1 to CE2

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



BGP Extensions for Administrative Flags

- CP level control flags: Policy Candidate Path Administrative Flags TLV
- Segment list control flags: Policy Segment list Administrative Flags TLV

```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
Attributes:
  Tunnel Encaps Attribute (23)
    Tunnel Type: SR Policy
    Binding SID
    SRv6 Binding SID
    Preference
    Priority
    Policy Name
    Policy Candidate Path Name
    Policy Candidate Path Administrative Flags
    Explicit NULL Label Policy (ENLP)
    Segment List
      Weight
      Segment List Administrative Flags
      Segment
      Segment
      ...
    ...
```

BGP Extensions for CP Flags

Add SR CP Admin Flags TLV, Type: TBD1



```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
Attributes:
  Tunnel Encaps Attribute (23)
    Tunnel Type: SR Policy
    Binding SID
    SRv6 Binding SID
    Preference
    Priority
    Policy Name
    Policy Candidate Path Name
    Policy Candidate Path Administrative Flags
    Explicit NULL Label Policy (ENLP)
    Segment List
      Weight
      Segment List Administrative Flags
      Segment
      Segment
      ...
```

S-Flag: When set, indicates that this CP is in admin Shut state.

B-Flag: When set, Indicates the CP is configured as "backup ineligible"

I-Flag: Indicates the CP is configured as "Ignore service route's Prefix SID".

D-Flag: Indicates the CP has been marked as the "Drop Upon Invalid" behavior as described in section 8.2 of [RFC9256].

BGP Extensions for Segment list Flags

Added SR Segment List Admin Flags TLV, carried within the Sub-TLVs of the SR Segment List TLV, Type: TBD2


0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length										Flags																			

0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6
S E																

S-Flag: When set, indicates that this Seglist is in admin Shut state.

```

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
Attributes:
  Tunnel Encaps Attribute (23)
    Tunnel Type: SR Policy
    Binding SID
    SRv6 Binding SID
    Preference
    Priority
    Policy Name
    Policy Candidate Path Name
    Policy Candidate Path Administrative Flags
    Explicit NULL Label Policy (ENLP)
    Segment List
      Weight
      Segment List Administrative Flags
      Segment
      Segment
      ...
      ...
  
```



Next Steps

- Any questions or comments are Welcomed
- Seeking for feedback

THANKS