

# **BGP SR Policy Extension for Metric**

<https://datatracker.ietf.org/doc/draft-zhang-idr-sr-policy-metric/>

K. Zhang (Huawei Technologies)

J. Dong (Huawei Technologies)

# Background

- BGP can be used to propagate the SR Policy candidate paths to the headend nodes in the network.
- The SR Policy Encoding structure is as follows:

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

Segment List

Weight

Segment

...

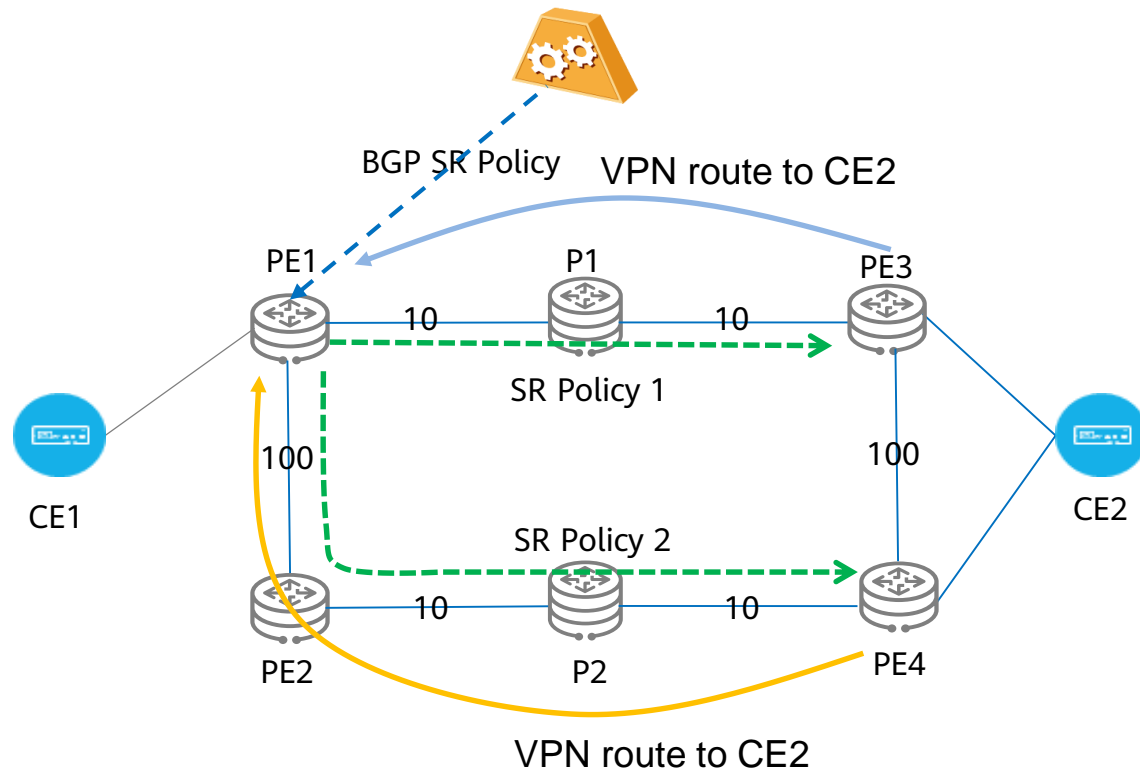
...

# Motivation & Problem Statement

By using BGP SR Policy NLRI, controllers may deliver different SR Policies which satisfy the constraints of VPN services. In some cases, the VPN route need choose next hop according to the IGP metric of each SR Policy.

For example, in this scenario:

- on PE1, the VPN routes to CE2 has two different next hop PE3 and PE4, each has one SR Policy: SR Policy1 and SR Policy2.
- Both SR Policies satisfies the delay and bandwidth constraints. And the IGP metric of the actual path of SR Policy1, which is PE1-P1-PE3, is 20; whereas the IGP metric of the actual path of SR Policy2, which is PE1-PE2-P2-PE4, is 120.
- In this case, PE1 wants to choose the next hop according to the IGP metric of SR Policies, thus can have the path with less cost.



Thus controller needs to deliver the IGP metric of SR Policies using BGP SR Policy.

# BGP SR Policy Extensions

- To solve this problem, this document proposes to add the Metric Sub-TLV to BGP SR Policy. The Metric Sub-TLV is a Sub-TLV of Segment List.

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

Explicit NULL Label Policy (ENLP)

Segment List

Weight

**Metric**

Segment

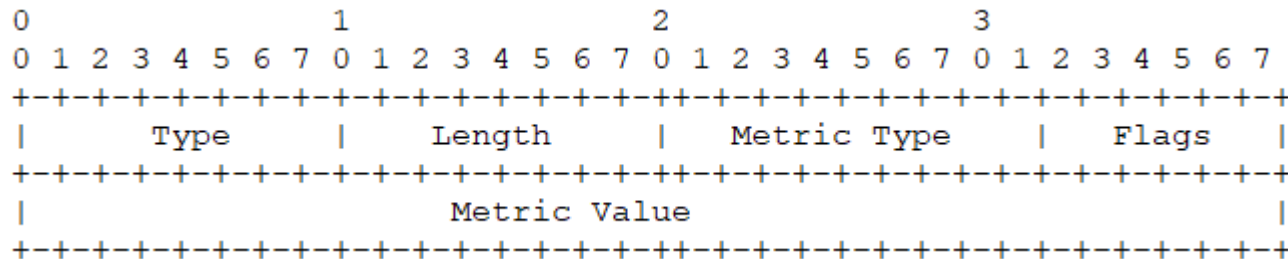
Segment

...

...

# BGP SR Policy Extensions

The format of Metric Sub-TLV is defined as follows:

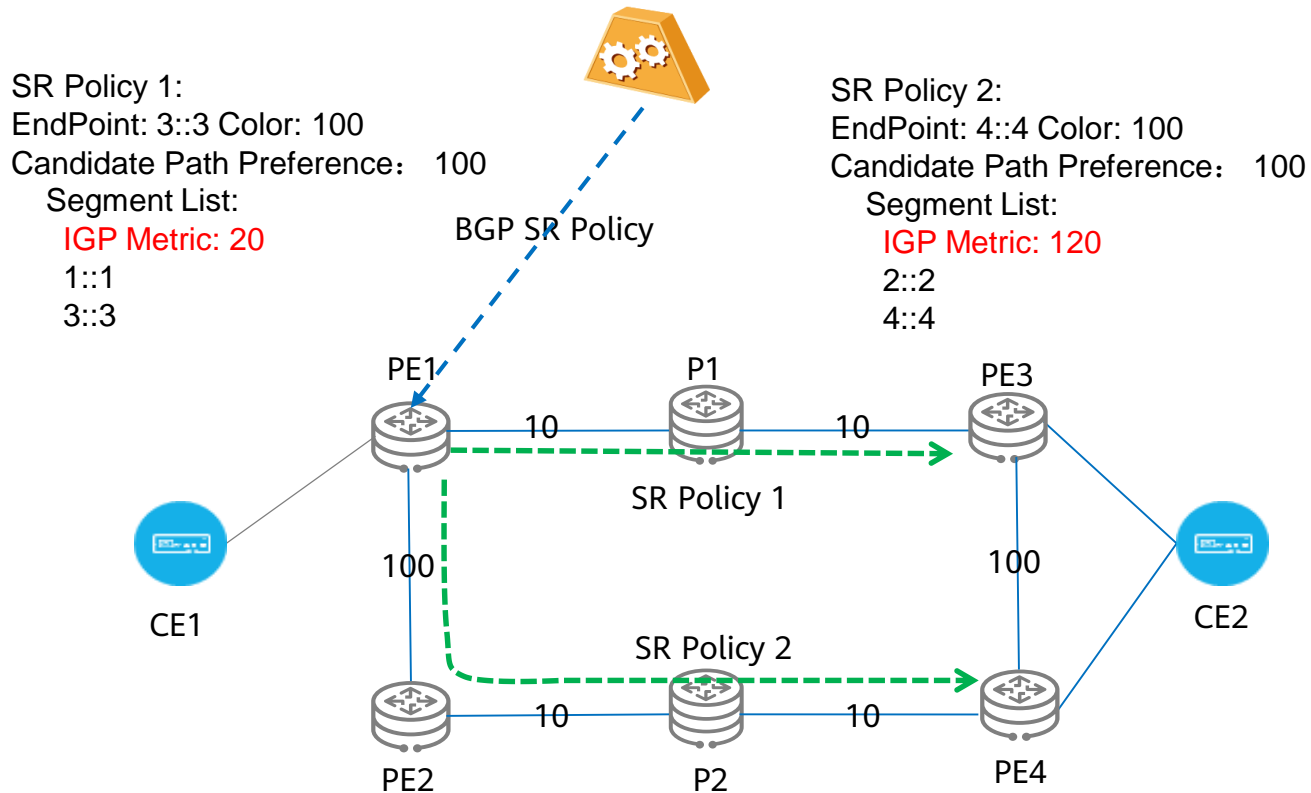


Where:

- Type: 1 octet, TBD.
- Length: 6 octets.
- Flags: 1 octet of flags. None are defined at this stage. Flags SHOULD be set to zero on transmission and MUST be ignored on receipt.
- Metric Type: 1-octet field which identifies the type of the metric being used.
- Metric Value(32 bits): 4-octet value which indicates the metric of the computed path.

Code Point	Metric Type
0	IGP Metric
1	Min Unidirectional Link Delay [RFC7471]
2	TE Metric [RFC3630]
3	Hop Count (refer [RFC5440])

# Processing of the Metric



- PE1 has two different SR Policies.
  - The next hop to PE3 uses SR Policy1, metric is 20.
  - The next hop to PE4 uses SR Policy2, metric is 120.
- PE1 chooses next hop according to IGP metric, and the result next hop is PE3.

# Processing of the Metric (cont.)

When SR Policy headend gets the SR Policy segment list with metric, how to process the metric is local policy.

- The active candidate path of SR Policy may have several segment lists, each segment list may have different metric.
- It is recommended that the segment lists in one candidate path have the same metric type.
- If the metric values of segment lists in one candidate path are different, the candidate path can use the maximum value as the metric of the candidate path.
- And the SR Policy use the metric value of the active candidate path.

The SR Policy metric value then can be used in routing preference.

# Updates

- There are some updates in the latest version (-03).
  1. Add the min unidirectional link delay metric Type.
  2. Add codepoints for different metric types.
  3. Add text about the candidate path metric calculation
    - When calculating the metric of the candidate path, if the metric of the segment lists are different, the maximum metric value of the segment lists is used as the metric of the candidate path



# Next Steps

- Welcome comments and questions
- WG adoption?

Thank you!