

IGP Color-Aware Shortcut

draft-cheng-lsr-igp-shortcut-enhancement-04

Presenter: Changwang Lin (New H3C Technologies)

Co-authors: Weiqiang Cheng (China Mobile)

Liyan Gong (China Mobile)

Changwang Lin (New H3C Technologies)

Mengxiao Chen (New H3C Technologies)


IETF-120, July 2024

Motivation

IGP shortcut [RFC3906] enables IGP to calculate routes over Traffic Engineering tunnels.

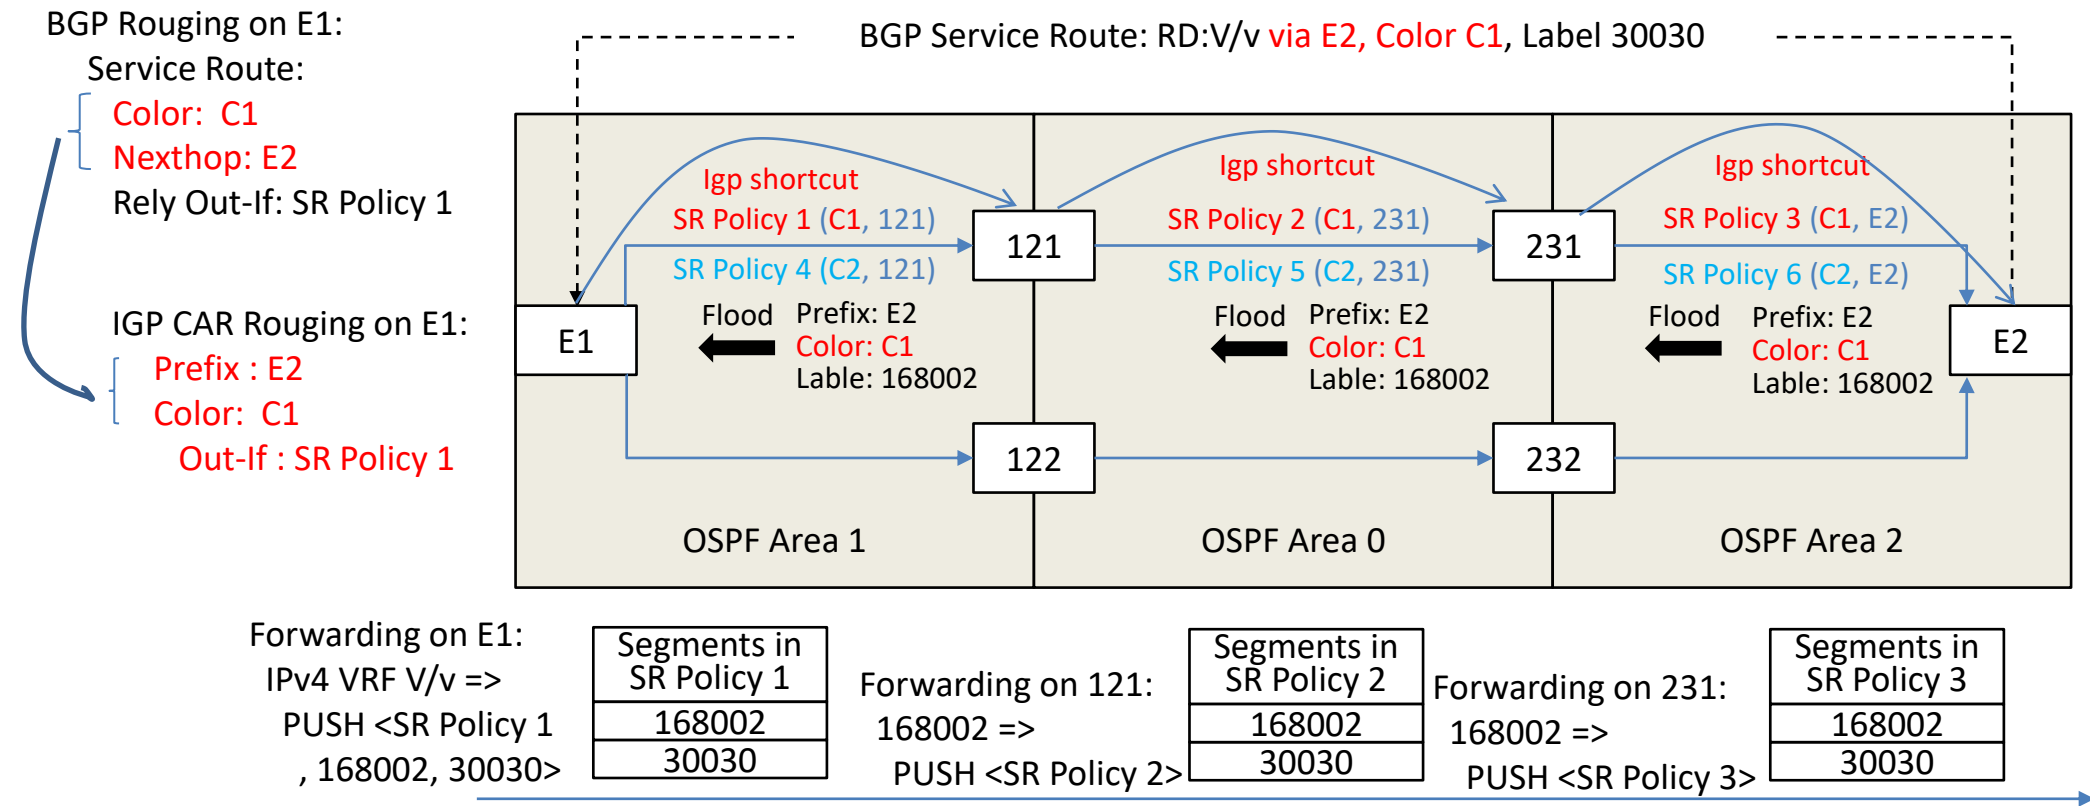
The granularity of IGP shortcut is based on nodes. If the first-hop of a node is determined to be a TE-tunnel during the SPF computation, all routes to IP prefixes advertised by that node will be over that TE-tunnel.

In some scenarios, there may be requirements to steer the routes to different prefixes of the same node onto different TE-tunnels. This document describes the enhancement of IGP shortcut which can steer routes onto TE-tunnels based on colors.



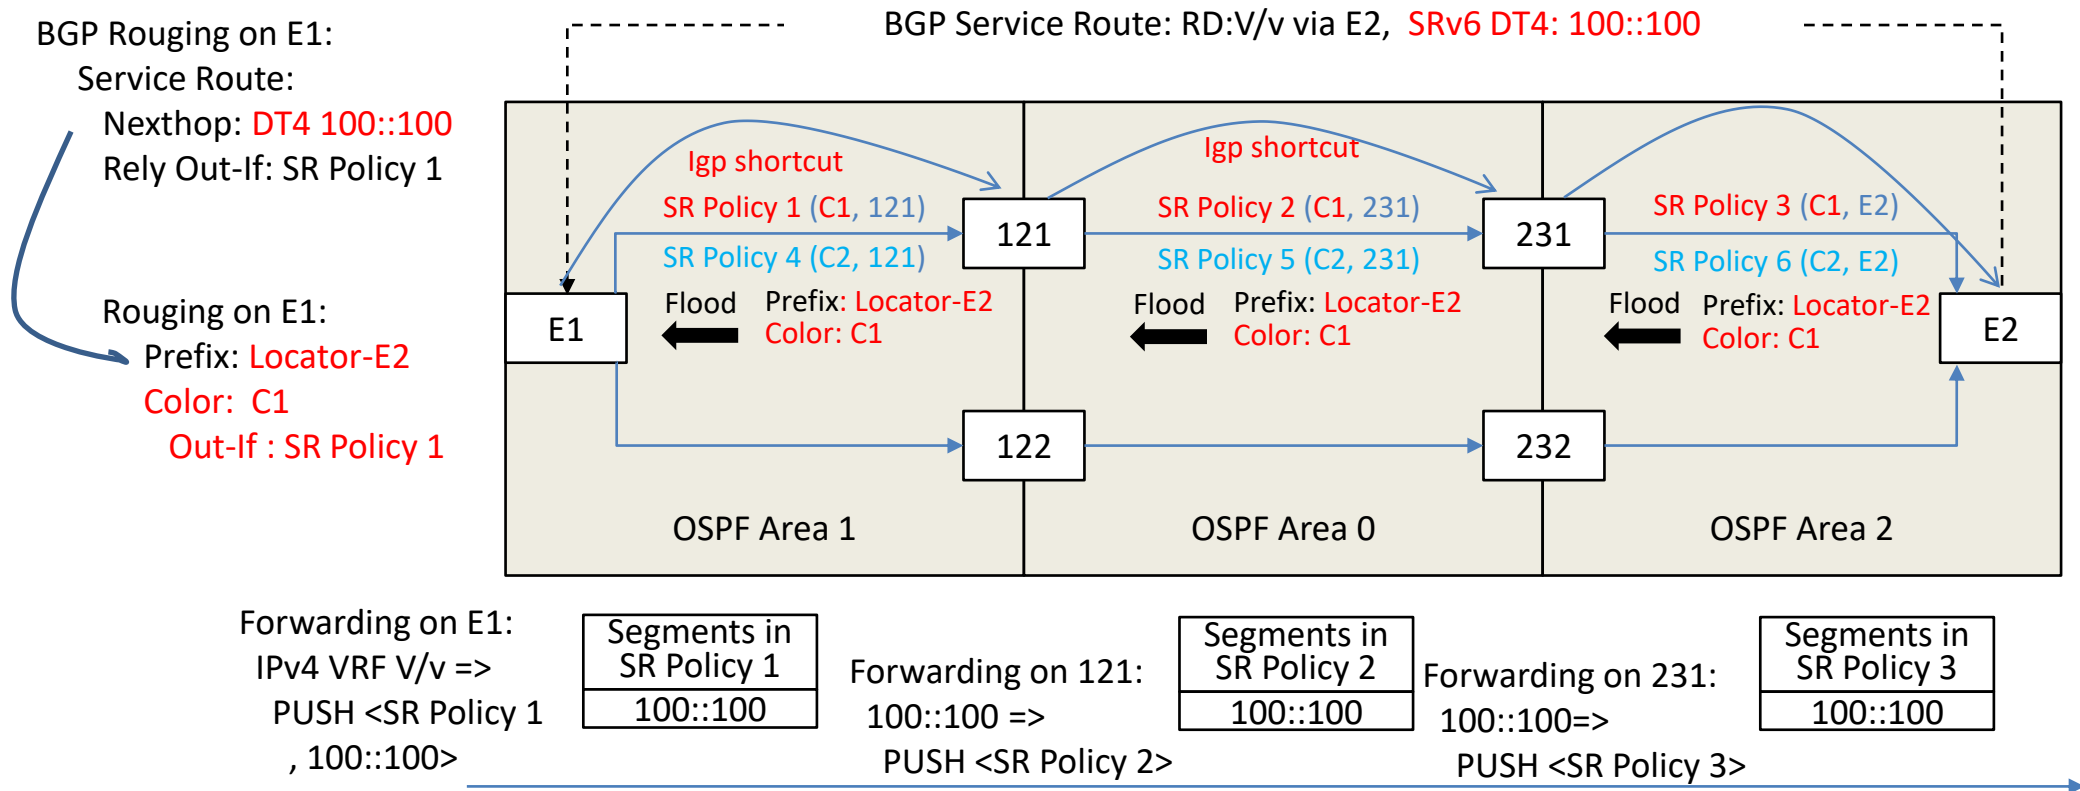
Use Case 1: SR Policy based IGP CAR

- SR Policy provides intent in each area.
- IGP Shortcut run in each area.
- IGP CAR route [draft-lin-lsr-igp-car] is resolved over color-matched SR Policy.



Use Case 2: SRv6 Locator based IGP CAR

- The IGP advertises Locator route in each area: Locator-E, Color=C1.
- Shortcut calculation is performed in each area, and the resulting route points to the SR Policy associated with the corresponding color.




Solution Overview

There is already implementation of policy-based shortcut. On the headend, filtering polices can be configured for tunnels used in IGP shortcut. Only the prefixes which satisfy the polices could be steered over the associated tunnel. However, policy-based shortcut may create difficulty for operation and maintenance: $N_{\text{config}} = N_{\text{headend}} * N_{\text{tunnel}} * N_{\text{prefix}}$

Proposed solution:

- Colored IGP prefix.
- Colored prefixes are steered over color-matched tunnels.

(Some kinds of TE-tunnels, such as SR Policy, have inherent color values. For the TE-tunnels which have no inherent color, the color values may be determined by local configurations.)



IGP Color-Aware Shortcut

Advertising Colors for IGP Prefixes

Approach #1: Using the existing Administrative Tag

Use IGP tags to colored prefixes:

- OSPF: 32-bit Administrative Tag TLV [draft-ietf-lsr-ospf-admin-tags]
- IS-IS: 32-bit Administrative Tag Sub-TLV [RFC5130]

Configure the mapping between tags and colors.

- Advantage: No new sub-TLV needs to be defined.
- Disadvantage: Multiple tags could be attached to one prefix. There must be a mechanism for the receiver to know which tag is used as color.

Approach #2: Defining a new Color Sub-TLV

The IS-IS Color Sub-TLV is applicable to TLVs 27, 135, 235, 236, and 237.

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	0	1
Type										Length										Flags																					
Color																																									

The OSPF Color Sub-TLV is applicable to OSPFv2 Extended Prefix TLV、OSPFv3 Inter-Area-Prefix TLV、OSPFv3 Intra-Area-Prefix TLV、OSPFv3 External-Prefix TLV.

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	0	1
Type										Length										Flags										Color											
Color (cont.)																																									

Multiple Colors Sub-TLVs could be attached to one prefix.

Update SPF Computation [RFC3906]

Update the SPF computation of IGP shortcut [RFC3906]:

- Examine the list of tail-end routers directly reachable via a TE-tunnel. If there is a TE-tunnel to this node, we use the TE-tunnel as the first-hop.
- If there is no TE-tunnel, and the node is directly connected, we use the first-hop information from the adjacency database.
- If the node is not directly connected, and is not directly reachable via a TE-tunnel, we copy the first-hop information from the parent node(s) to the new node.



If there is a TE-tunnel to this node, we **add** the TE-tunnel into the first-hop information (**without deleting the previous ones**).

Adds the following steps when calculating next-hops for prefixes advertised by a node:

If a prefix is colored, we look up the first-hop information of the advertiser node for TE-tunnels with the same color.

- If there are eligible TE-tunnels, we compare the costs of paths over those TE-tunnels, and use the next-hop of the TE-tunnel with the lowest path cost.
- If there is no eligible TE-tunnel, we use the native adjacency next-hop.

If a prefix has no color, we use the next-hop with the lowest path cost.

Updates and Status

- Update to use tags for color mapping.
- Include use cases for SR-MPLS and SRv6.

Next Steps

- Any questions or comments are welcomed.
- Seeking for adoption.

Thanks