

IGP Color-Aware Routing

draft-lin-lsr-igp-car-00

Changwang Lin (New H3C Technologies)

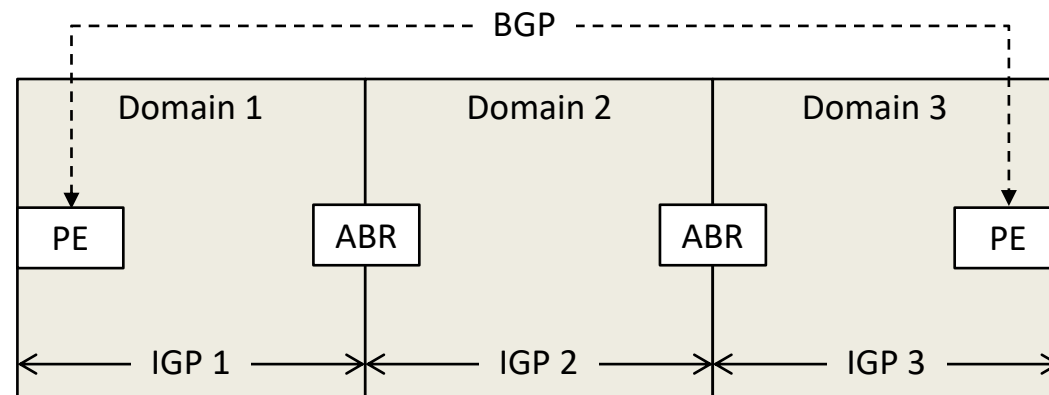
Mengxiao Chen (New H3C Technologies)

Liyan Gong (China Mobile)

IETF-118, November 2023

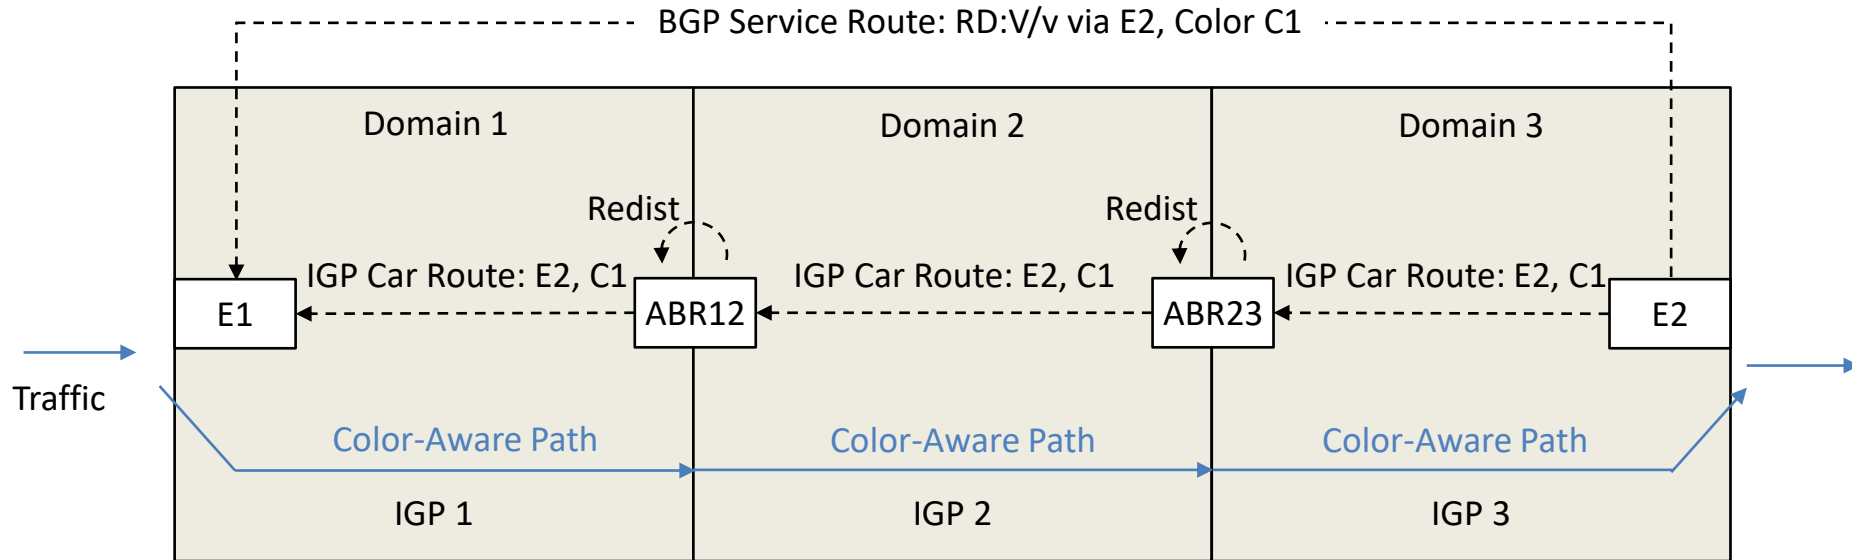
Background

- There are requirements for end-to-end intent-aware paths spanning multiple domains, as described in [I-D.hr-spring-intentaware-routing-using-color].
- In some networks, only the PE nodes run BGP, while the other nodes run only IGP.



- Existing works focus on BGP-based solutions: BGP-CAR [I-D.ietf-idr-bgp-car] & BGP-CT [I-D.ietf-idr-bgp-ct]. They require the ABR nodes to run BGP and signal BGP routes for CAR path.
- This draft proposes an IGP-based solution without requiring the ABR nodes to run BGP.

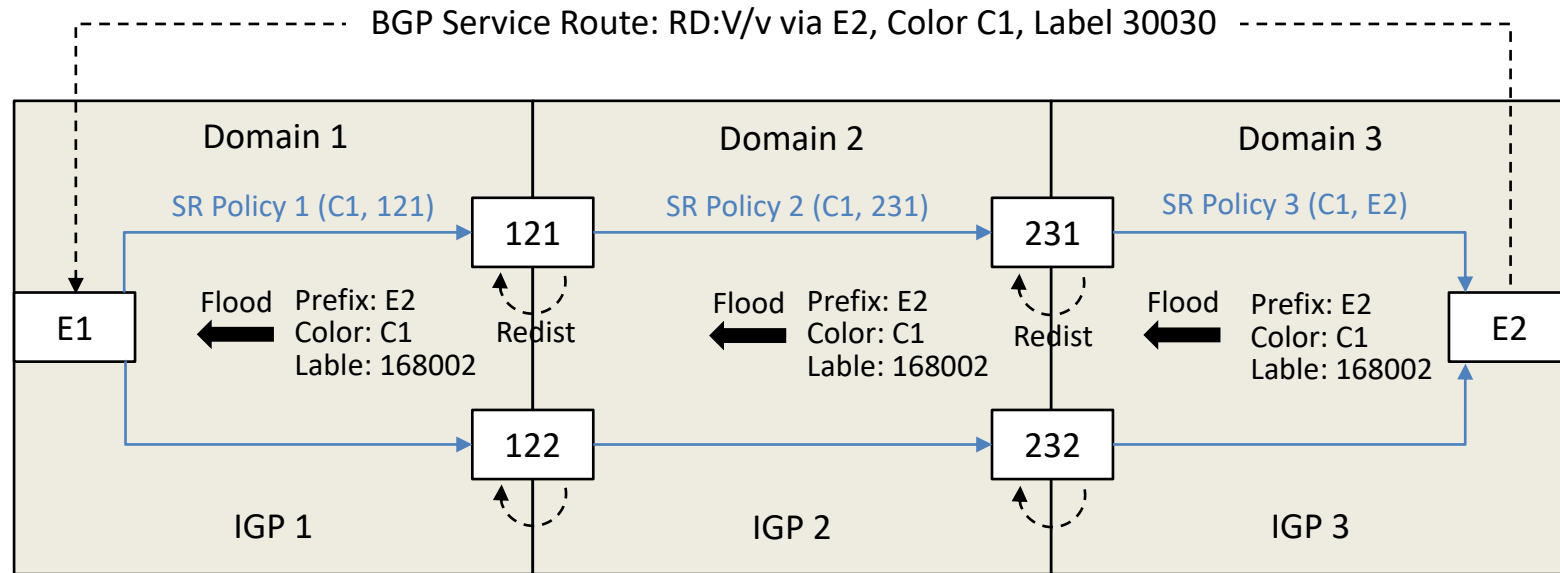
Solution Overview



- A BGP service route RD:V/v is colored with C1 using BGP Color Extended-Community.
- (E2, C1) is an IGP CAR route in underlay that provides intent-aware path to E2. It is originally advertised by E2, and redistributed by ABR nodes across different IGP domains.
- BGP service route is resolved over IGP CAR route, and automatically steered onto a Color-Aware path.
- Steering for all services (L3VPN, EVPN, Internet/global table etc).

SR Policy based IGP CAR

- SR Policy provides intent in each domain.
- On each ABR and the Ingress PE E1, CAR label 168002 is resolved over SR Policy of the domain.



Forwarding on E1:
IPv4 VRF V/v =>
PUSH <SR Policy 1
, 168002, 30030>

Segments in SR Policy 1
168002
30030

Forwarding on 121:
168002 =>
PUSH <SR Policy 2>

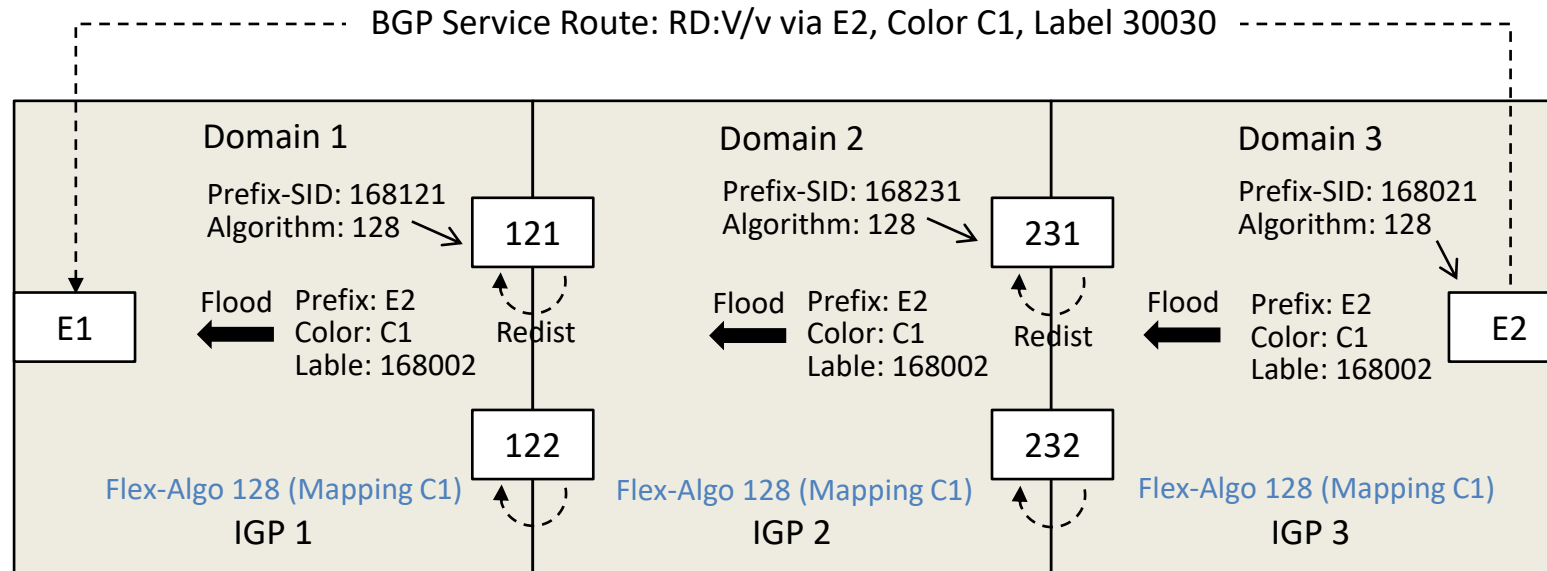
Segments in SR Policy 2
168002
30030

Forwarding on 231:
168002 =>
PUSH <SR Policy 3>

Segments in SR Policy 3
168002
30030

Flex-Algo based IGP CAR

- Flex-Algo 128 is running in each domain, and mapped to Color C1.
- On each ABR and the Ingress PE E1, CAR label 168002 is resolved over IGP Flex-Algo 128 of the domain.



Forwarding on E1:

IPv4 VRF V/v =>
 PUSH <168121
 , 168002, 30030>

168121
168002
30030

Forwarding on 121:
 168002 =>
 PUSH <168231>

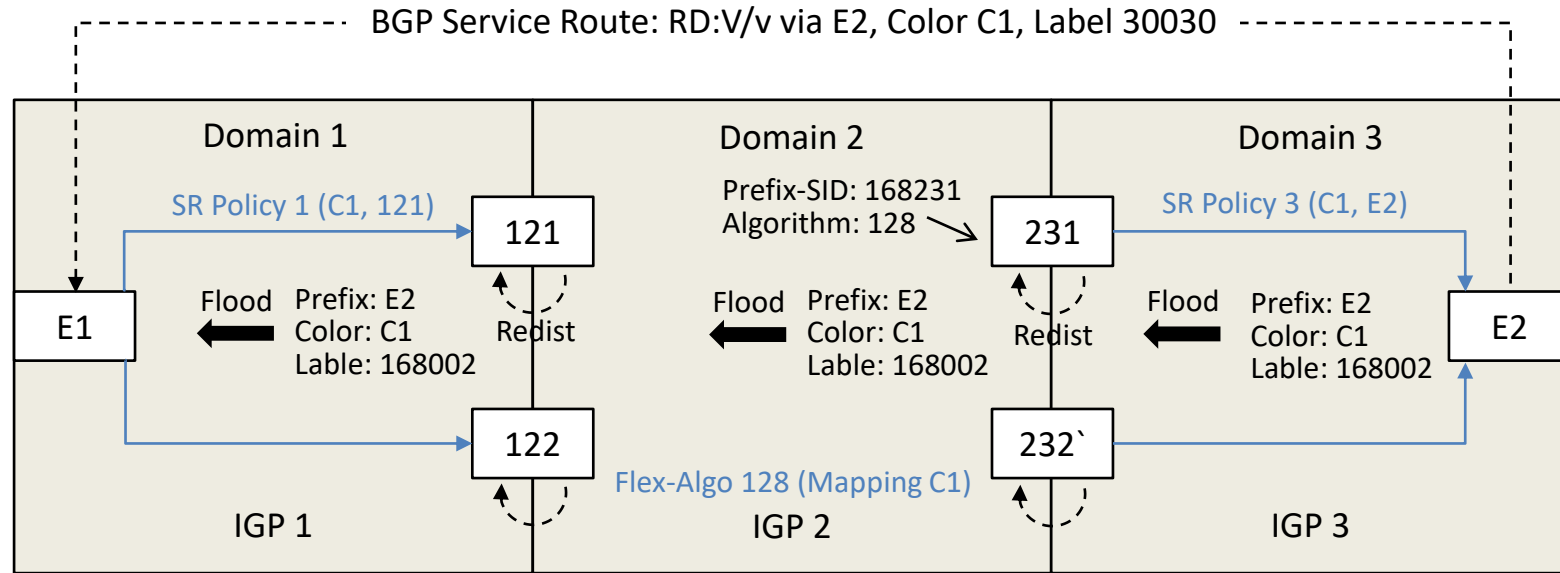
168231
168002
30030

Forwarding on 231:
 168002 =>
 PUSH <168021>

168021
168002
30030

Hybrid IGP CAR

- In domain 1 and 3, the intent of color C1 is provided by SR Policy.
- In domain 2, C1 is mapped to Flex-Algo 128.



Forwarding on E1:
IPv4 VRF V/v =>
PUSH <SR Policy 1
, 168002, 30030>

Segments in SR Policy 1
168002
30030

Forwarding on 121:
168002 =>
PUSH <168231>

168231
168002
30030

Forwarding on 231:
168002 =>
PUSH <SR Policy 3>

Segments in SR Policy 3
168002
30030

Advertisement of CAR Route in IGP

- The IS-IS CAR Sub-TLV is used to advertise CAR information for prefixes in IS-IS. It is applicable to TLVs 27, 135, 235, 236, and 237.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|   Type   |   Length   |   Flags   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Color                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|Sub-sub-TLV-len|   Sub-sub-TLVs (variable)...   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Carrying the encapsulation information:

- MPLS label/index
- SRv6 SID

- The OSPF CAR Sub-TLV is used to advertise CAR information for prefixes in OSPFv2 and OSPFv3. It is applicable to OSPFv2 Extended Prefix TLV, OSPFv3 Inter-Area-Prefix TLV, OSPFv3 Intra-Area-Prefix TLV, and OSPFv3 External-Prefix TLV.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           Type           |           Length           |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           Flags           |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Color                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Sub-sub-TLVs (variable)...   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Example:

IGP Prefix TLV: Prefix-1

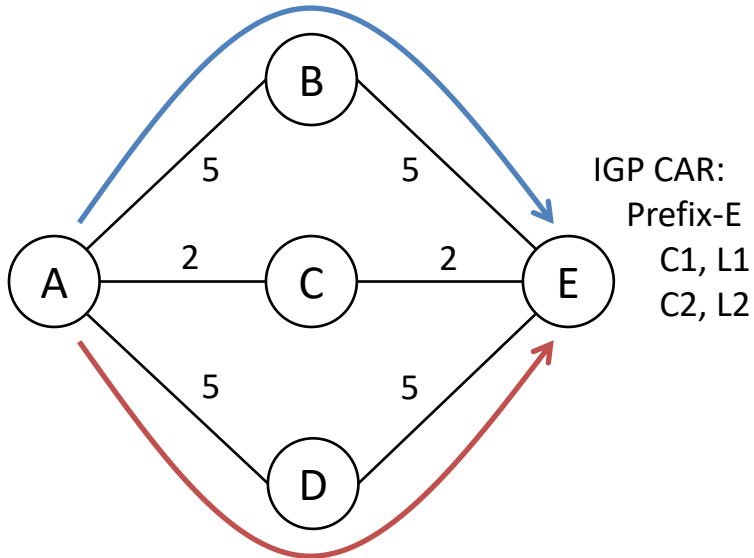
CAR Sub-TLV: Color-1 + Label-1

CAR Sub-TLV: Color-2 + Label-2

CAR Sub-TLV: Color-3 + Label-3

Resolving of IGP CAR over SR Policy

SR Policy 1 (C1, E), Seglist=S1, cost=1



SR Policy 2 (C2, E), Seglist=S2, cost=1

IGP Short-cut (RFC3906):

- E is resolved over SR Policy 1 & 2 with load balancing

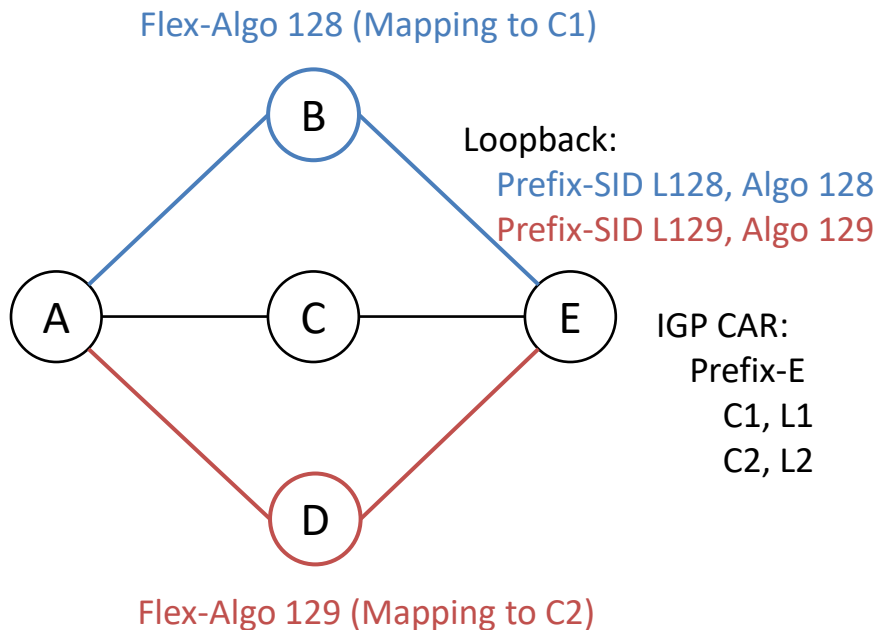
Color-aware IGP Short-cut (draft-cheng-lsr-igp-shortcut-enhancement):

- (E, C1) is resolved over SR Policy 1 ✓
- (E, C2) is resolved over SR Policy 2 ✓

MPLS Forwarding Entries for IGP CAR Route:

In-Label	Out-Label	Next-Hop
L1	<S1, L1>	B (Along the path of SR Policy 1)
L2	<S2, L2>	D (Along the path of SR Policy 2)

Resolving of IGP CAR over Flex-Algo



- ① Determine the FA to which the color is mapped:
 - (E, C1) is resolved over FA 128
 - (E, C2) is resolved over FA 129
- ② Add the CAR advertiser's Prefix-SID associated with that FA into encapsulation.
 - (E, C1) is encapsulated by L128
 - (E, C2) is encapsulated by L129

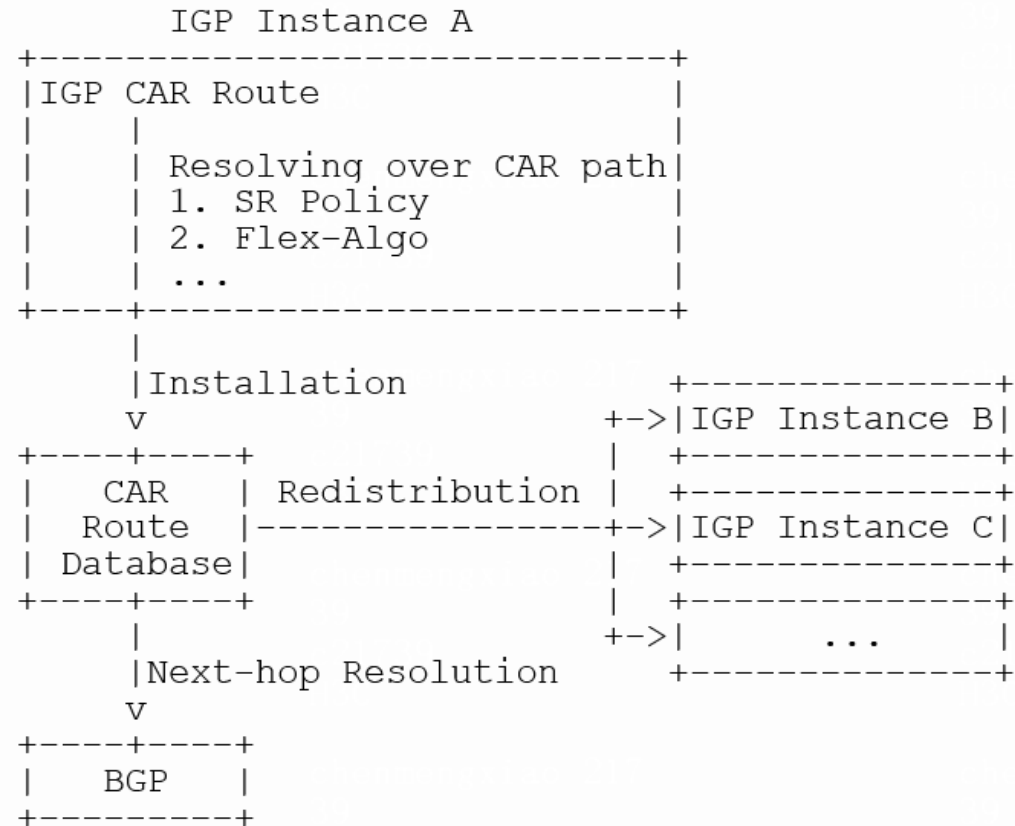
MPLS Forwarding Entries for IGP CAR Route:

In-Label	Out-Label	Next-Hop
L1	<L128, L1>	B (Along the path in Flex-Algo 128)
L2	<L129, L2>	D (Along the path in Flex-Algo 129)

CAR Route Database (CARDB)

CARDB is a logical collection of resolved CAR routes:

- When the ABR or Ingress PE receives the IGP CAR route, it will be resolved over the CAR paths, and then installed to CARDB.
- The ABR can redistribute the CAR routes in CARDB to other IGP instances.
- The Ingress PE can use the CARDB for next-hop resolution of BGP colored service route.



Comparisons with BGP-based Solution

Pros:

- Do not require the ABR nodes to run BGP. No need to change existing routing protocol deployments, when enabling intent-aware routing in IGP-only networks.

Cons:

- CAR routes (E, C) are flooded to all IGP nodes. (Opaque to intermediate nodes. Only ABR nodes have to deal with CAR information. Summarization of CAR routes might help.) Suitable for not very large networks only.
- Color assignments of different domain must be unified.

Next Steps

- Add details about SRv6 data plane and summarization of CAR route.
- Any questions or comments are Welcomed.

Thanks